

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**



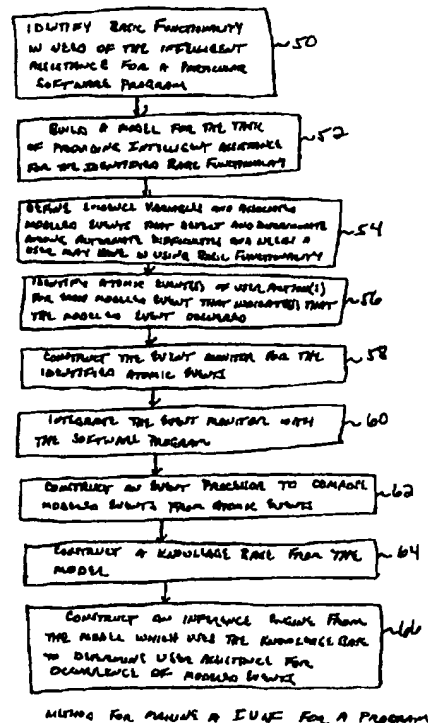
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/00		A2	(11) International Publication Number: WO 98/03907
			(43) International Publication Date: 29 January 1998 (29.01.98)
(21) International Application Number: PCT/US97/10856		(81) Designated States: CN, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 2 July 1997 (02.07.97)		Published Without international search report and to be republished upon receipt of that report.	
(30) Priority Data: 08/684,003 19 July 1996 (19.07.96) US			
(71) Applicant: MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, WA 98052-6399 (US).			
(72) Inventors: HORVITZ, Eric; 330 Waverly Way, Kirkland, WA 98033 (US). BREESE, John, S.; 7701 S.E. 58th Street, Mercer Island, WA 98040 (US). HECKERMAN, David, E.; 648 West Lake Sammamish Lane N.E., Bellevue, WA 98008 (US). HOBSON, Samuel, D.; 98 Union Street #905, Seattle, WA 98101 (US). HOVEL, David, O.; 4610 East Lake Sammamish Parkway N.E., Redmond, WA 98053 (US). KLEIN, Adrian, C.; 411 Boylston Avenue E. #203, Seattle, WA 98102 (US). ROMMELSE, Jacobus, A.; Braambos 20, NL-5563 AB Westervoort (NL). SHAW, Gregory, L.; 5402 Lake Washington Boulevard N.E. #E, Kirkland, WA 98033 (US).			
(74) Agents: CANNING, Kevin, J. et al.; Seed and Berry LLP, 6300 Columbia Center, 701 Fifth Avenue, Seattle, WA 98104-7092 (US).			

(54) Title: INTELLIGENT USER ASSISTANCE FACILITY

(57) Abstract

A general event composing and monitoring system that allows high-level events to be created from combinations of low-level events. An event specification tool allows for rapid development of a general event processor that creates high-level events from combinations of user actions. The event system, in combination with a reasoning system, is able to monitor and perform inference about several classes of events for a variety of purposes. The various classes of events include the current context, the state of key data structures in a program, general sequences of user inputs, including actions with a mouse-controlled cursor while interacting with a graphical user interface, words typed in free-text queries for assistance, visual information about users, such as gaze and gesture information, and speech information. Additionally, a method is provided for building an intelligent user interface system by constructing a reasoning model to compute the probability of alternative user's intentions, goals, or informational needs through analysis of information about a user's actions, program state, and words. The intelligent user interface system monitors user interaction with a software application and applies probabilistic reasoning to sense that the user may need assistance in using a particular feature or to accomplish a specific task. The intelligent user interface also accepts a free-text query from the user asking for help and combines the inference analysis of user actions and program state with an inference analysis of the free-text query. The inference system accesses a rich, updatable user profile system to continually check for competencies and changes assistance that is given based on user competence.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
ER	Estonia						

Description

INTELLIGENT USER ASSISTANCE FACILITY

5 Technical Field

 This invention relates to expert systems in data processing. In particular, this invention relates to a system and method for automated machine reasoning to provide assistance to software users and to optimize the functionality of computer systems and software by performing inference about users needs and preferences in the
10 operation of software systems or applications.

Background of the Invention

 Although significant advances in graphical user interfaces have made software applications more accessible and productive for personal computer users, the
15 increased ease of use has fostered a demand to include more and more sophisticated features. When first exposed to a complicated software application, a new user can be confused by the overwhelming number and complexity of features accessible through the menus and toolbars of the user interface. It is not uncommon for a user unfamiliar with the software features to resort to menu surfing, rapidly switching from menu item
20 to menu item, in the hope of discovering how to effect the desired feature. Even when users know one or more ways to perform a task in a software application, they may not realize an efficient approach. Also, user's may know how to perform tasks, yet instead of performing the task themselves, they would rather have a system proactively determine their needs and perform or offer to perform operations, such as launching
25 applications that will soon be needed, exiting applications, and prefetching files or information from distant servers to make them available more efficiently when they are requested.

 Determining in an automated manner the best actions to perform or the best information to provide users while they work with software requires the
30 development of computational methods that operate in conjunction with software

programs and that have the ability to identify a user's needs, intentions, or goals from aspects of a user's background and actions. Such methods may be derived from the construction and use of models that explicitly links the users' needs to their backgrounds and actions.

5 There have been several studies of the use of models for reasoning about the intentions of people performing a variety of tasks, and for making available advice or assistance based on these models. Researchers studying the use of computers in education have attempted to construct programs that use models of users to determine the source of misunderstanding and the best way to tutor a student. These models can
10 look at the answers users give to questions and other challenges. The Strategic Computing Initiative of the 1980s focused in part on methods for doing reasoning about a pilot's intentions from a diverse set of inputs, and for providing them with the best assistance with piloting and warfighting with fighter jets. Within the many such "pilot associates" projects, engineers attempted to design systems that could reason about the
15 intentions of users. Typically, rule-based or logical approaches were used for determining or prioritizing options.

 There has also been some work in the use of probabilistic models in the modeling of user's needs for such tasks as monitoring complex, time-critical applications. Some of the earliest discussion of probabilistic dependency models, with
20 applications to enhancing the display of information, were in projects experimenting with the modeling of users for controlling the information displayed to aircraft pilots. Probabilistic inference was explored as a means for controlling the information being displayed to pilots. Other work included using decision-theoretic methods to reason about the most important information to display to flight engineers at NASA Mission
25 Control Center. The methods consider context, user expertise level, and telemetry from the Space Shuttle to make these decisions, and take advantage of time-criticality and probabilistic models of user beliefs and intentions. Other related work has explored the best kind of information and explanations to display to pilots given various tradeoffs in display. In other work, researchers have explored the use of probabilistic models in

assisting computer users working with command line input systems of a computer operating system.

There have been several projects focusing on the use of sets of logical rules and heuristics based on logical rules for providing users with assistance based on the context and activity. For instance, the common context-sensitive help that is made available in applications written for the Microsoft Windows operating system is accessed by selecting a graphical object and clicking a standard help command (e.g., a function key on the keyboard). Help information is then provided which describes the functionality of the graphical object such a toolbar. Another example is the TipWizard system in Microsoft Excel products which employs a large set of logical rules that continue to scan for sequences of user events and provide assistance to users on performing tasks more efficiently. In work at Apple Computer described by U.S. Patent No. 5,390,281 to Luciw et al., sets of rules are used for providing assistance to users on tasks by employing the heuristic of comparing the number of conditions for a rule to fire with the specific number of conditions that are seen in a particular situation and for interpreting this fraction as a strength associated with the action indicated by the rule.

There has also been work on more sophisticated probabilistic models that consider the uncertainty in user goals and needs for software assistance. The Answer Wizard feature of Microsoft Office products on the use of probabilistic models for interpreting the words in a user's free-text query with a focus on assisting users with using computer software applications. In this work, the uncertain relationship between sets of related words that might be seen in a query and various kinds of assistance are modeled with probability. This work has been described in U.S. Patent Application No. 08400797.

However, none of these former efforts have employed sophisticated methods for reasoning under uncertainty about a user's needs for assistance by considering multiple pieces of evidence, including information about a user's background and a user's activities that are autonomously sensed by the program, and, then, have offered to provide relevant assistance to a user based on this reasoning. Nor have these systems considered monitoring user interaction with a software program so

that, when a user does inquire about help directly with free-text queries, the system combines the analysis of the words in the free-text query with the ongoing analysis of user actions and program context.

What is desirable is an intelligent user assistance facility that autonomously senses that the user may need assistance in using a particular feature or to accomplish a specific task, and that offers to provide relevant assistance based upon considering multiple pieces of evidence, including information about a user's background and a user's activities. What is also desirable is an intelligent user assistance facility that combines the analysis of the words in the free-text query with the ongoing analysis of user actions and program context whenever a user asks directly for help with free-text queries. This invention solves the problem of autonomously sensing that a user may need assistance in using a particular feature or to accomplish a specific task, and offering to provide relevant assistance based upon considering multiple pieces of evidence, including information about a user's background and a user's activities. This invention also solves the problem of combining the analysis of the words in the free-text query with the ongoing analysis of user actions and program context whenever a user asks directly for help with free-text queries. In addition, the invention is able to monitor and perform inference about several classes of events, including the state of key data structures in a program, general sequences of user inputs with a mouse-controlled cursor in the normal course of interacting with a graphical user interface, words typed in free-text queries for assistance, visual information about users, such as gaze and gesture information, and speech information.

Summary of Invention

The present invention provides a method of building an intelligent user assistance facility for a software program to assist a user in operation of the software program. Functionality of the software program is identified for providing intelligent user assistance. A reasoning model is constructed for the identified program functionality for the task of providing intelligent user assistance. The reasoning model computes the probability of alternative user's intentions, goals, or informational needs

through analysis of information about a user's actions, program state, and words. Modeled events for the identified program functionality are defined for use by the reasoning model to provide intelligent user assistance. Atomic user interface interactions are identified that signify the occurrence of modeled events.

5 The present invention also includes a general event composing and monitoring system that allows high-level events to be created from combinations of low-level events. The event system, in combination with a reasoning system, is able to monitor and perform inference about several classes of events for a variety of purposes. The various classes of events include the state of key data structures in a program,
10 general sequences of user inputs with a mouse-controlled cursor in the normal course of interacting with a graphical user interface, words typed in free-text queries for assistance, visual information about users, such as gaze and gesture information, and speech information. The present invention additionally provides a system and method to monitor user interaction with a software program and to apply probabilistic reasoning
15 to sense that the user may need assistance in using a particular feature or to accomplish a specific task. The system continues to examine a user's actions and, from this information, generates probabilities about user needs and goals.

 In an exemplary embodiment, the likelihoods that various text-based help topics or demonstrations of various software tasks are relevant are computed and
20 the assistance is made available to users. Probabilities are computed about relevant help topics as well as about the probability that a user would desire help. In this embodiment, the invention computes the probabilities that a user needs alternate forms of help, and uses this information to prioritize a list of help topics, when help is requested. When help has not been requested, the system uses its inference to
25 determine when a user might desire assistance, and comes forward autonomously, depending on the computed importance of coming forward and the level at which a user sets a threshold. Whenever a user asks directly for help with free-text queries, this embodiment combines the analysis of the words in the free-text query with the ongoing analysis of user actions and program context.

The invention includes a rich, updatable user profiling system. In the help-topic embodiment, the inference system accesses the user profile system to continually check for competencies and changes assistance that is given based on user competence.

5 The present invention also includes a novel inference system with an annotated Bayesian network to include special temporal reasoning procedures. Temporal reasoning mechanisms are used to consider the changes in relevance of events with current needs and goals of users with time passing after an event has occurred. The inference engine includes a new specific approximation procedure which uses a
10 single explicit Bayesian network knowledge base, but changes the likelihood information in the network based on the distance in the past that an observation was made.

 The system also allows users to control the thresholds that must be reached before users are offered assistance, allowing them to tailor the system's
15 behavior to their own personal preference about being distracted with assistance. Finally, the methods have been integrated with multiple graphical user interfaces, including a social-user interface in the form of an animated character.

Brief Description of the Drawings

20 Figure 1 is a block diagram of a computer system.

 Figure 2 is a flowchart for building an intelligent user assistance facility for a software program.

 Figure 3 is a diagram illustrating a simple two variable Bayesian network.

25 Figure 4 is a block diagram of an Intelligent User Assistance Facility.

 Figure 5 is a block diagram of an instrumented application program.

 Figure 6 is a block diagram of an instrumented spreadsheet application program for use with the Intelligent User Assistance Facility.

 Figure 7 is a block diagram of an Event Processor.

Figure 8 is a block diagram of an Event System Specification Tool for creating a modeled event composer.

Figure 9 is a block diagram of a User Profile System.

Figure 10 is a flowchart of the Event Status Update Routine of the User
5 Profile System of Figure 9.

Figure 11 is a flowchart of the User Profile Access Routine of the User Profile System of Figure 9.

Figure 12 is a flowchart of the Customized Tutorial Routine of the User Profile System of Figure 9.

10 Figure 13 is a block diagram of the Inference System of the present invention.

Figure 14 is a Bayesian network diagram illustrating the overall influence relationship for any software program using the IUAF of the present invention to provide user assistance.

15 Figure 15 is a diagram illustrating a portion of a Bayesian network representing some probabilistic dependencies over time.

Figure 16 is a flowchart of the value-of-information procedure according to the principles of the invention for acquiring user information.

20 Figures 17 - 22 illustrate screen outputs of the IUAF made in accordance with the principles of the present invention in an exemplary embodiment of a spreadsheet application.

Figure 23 illustrates screen output of the IUAF according to the principles of the present invention in an exemplary embodiment of a spreadsheet application for displaying a summary of a custom-tailored tutorial.

25 Figure 24 is a portion of a Bayesian network for IUAF made in accordance with the principles of the present invention.

Figure 25 is a portion of a Bayesian network for a spreadsheet application illustrating modeled and profiled events for charting data of a spreadsheet.

Figure 26 is a Bayesian network annotated with temporal dynamics information made in accordance with the principles of the present invention in an exemplary embodiment of a spreadsheet application.

5 Figures 27 is a portion of a Bayesian network illustrating a model in accordance with the principles of the present invention for assisting with exiting in a multiple application setting.

Figures 28 - 30 is a portion of a Bayesian network illustrating reasoning for assisting with exiting in a multiple application setting.

10 Figure 31 is a diagram illustrating the control modes of the IUAF Controller.

Figure 32 is a method according to the present invention for combining inference analysis of user actions and program state with inference analysis of free-text query.

15 Figure 33 is a flowchart of the autonomous assistance procedure of the present invention.

Figure 34 is an illustration of a social interface providing intelligent user assistance for a spreadsheet program.

20 Figure 35 is a block diagram of the IUAF with speech recognition system showing use of inference results to enhance speech recognition.

Detailed Description of the Invention

The present invention provides a general event composing and monitoring system that allows high-level events to be created from combinations of low-level events. The event system, in combination with a reasoning system, is able to
25 monitor and perform inference about several classes of events for a variety of purposes. The various classes of events include the state of key data structures in a program, general sequences of user inputs with a mouse-controlled cursor in the normal course of interacting with a graphical user interface, words typed in free-text queries for assistance, visual information about users, such as gaze and gesture information, and
30 speech information. The present invention also employs a novel event specification

tool that allows for rapid development of a general event processor that creates high-level events from combinations of user actions.

The present invention also includes a novel inference system with an annotated Bayesian network to include special temporal reasoning procedures. The inference engine includes a new specific approximation procedure which uses a single explicit Bayesian network knowledge base, but changes the likelihood information in the network based on the distance in the past that an observation was made. The knowledge base includes special annotations with temporal dynamics information and the inference engine includes special procedures to handle changes in relevance with the passing of time.

The present invention also provides an intelligent user assistance facility and method for use of probabilistic and logical reasoning to make decisions about the best way to provide assistance to a computer software user. The present invention additionally provides a method of building an intelligent user assistance facility by constructing a reasoning model to compute the probability of a user's intentions, goals, or informational needs through analysis of information about a user's actions, program state, and words. An exemplary embodiment of the present invention monitors user interaction with a software application and applies probabilistic reasoning to sense that the user may need assistance in using a particular feature or to accomplish a specific task. This embodiment of the present invention additionally accepts a free-text query from the user asking for help and combines the inference analysis of user actions and program state with an inference analysis of the free-text query. This embodiment of the present invention also incorporates a novel user profile system that stores information about the competence of users that may be established initially by user input or dialogue with the user and that is updated with the results of ongoing user activity. This persistent user profile provides user competency information during inference analysis and further enables customization of tutorial information for the user after an application session concludes.

Furthermore, this embodiment includes automated assistance reasoning to determine when a user might want help, and comes forward autonomously,

depending on the computed importance of coming forward and the level at which a user sets a threshold. When a user does inquire about help directly with free-text queries, the system also combines the analysis of the words in the free-text query with the ongoing analysis of user actions.

5 Figure 1 is a block diagram of a computer system 10 that is suitable for practicing an exemplary embodiment of the present invention. Those skilled in the art will appreciate that the computer of system 10 depicted in Figure 1 is intended to be merely illustrative and that the present invention may be practiced with other computer system configurations, including distributed systems and multiprocessor systems, and
10 handheld devices. The computer system 10 includes a processor 12 and an input/output system 14. The input devices may be, for example, a camera 15, a keyboard 16, a mouse 18, a microphone 20, a pointing device or other input device. The output devices may be, for example, a display screen 22, a printer 24, a speaker 26 or other output device. The computer system 10 may include a speech interface 27 that interfaces a
15 speech system with the microphone 20 and speaker 26. The computer system 10 may include a visual recognition and tracking interface 29 that interfaces a visual recognition and tracking system with the camera 15 or other specialized head gear with an infrared reflector or 3D ultrasound sensors. The computer system 10 may also include a network interface 28 that interfaces the computer with a network 30 that may be either a
20 local area network or a wide area network such as the Internet. The computer system additionally includes a system memory 32 and persistent storage 34 that contain application programs and the intelligent user assistance facility. The persistent storage may be a diskette, CDROM, hard-drive, or firmware. The computer system 10 includes an operating system 36 and several application programs such as a spreadsheet
25 application program 38, a word processing application program 40, or other application program 42. A system bus 44 interconnects the processor 12, system memory 32, persistent storage 34, and input/output system 14. Those skilled in the art will appreciate that the intelligent user assistance facility may be integrated directly into the application programs or may be a stand alone facility that is part of a system library or
30 operating system.

Overview of the Intelligent User Assistance Facility (IUAF)

Although software usability studies can pinpoint user interface problems for accessing complex features and can simplify and make the user interface more intuitive for accessing the features, improvements to the user interface cannot hide the multitude of complex features or simplify the functionality of a complex feature for the inexperienced user without reducing the number of features exposed through the interface and/or exposing only simplified versions of the full feature.

One by-product of software program usability studies can be the identification of user behavior when a user experiences difficulties in using basic functionality of a particular software program. This invention provides a process for making an intelligent user assistance facility for a particular software program by observing and understanding user behavior when a user experiences such difficulties or when a user might benefit by providing automated actions

Figure 2 is a flowchart of the method for making an intelligent user assistance facility for a software program. The first step 50 of the process is to identify the basic functionality of a particular software program in which users experience difficulties, or could benefit from automated assistance, and the user behavior exhibited by the user when the user is experiencing those difficulties or desire for assistance. One way of accomplishing this step is to observe and record user behavior when a user experiences difficulties during usability studies for that particular software program. For example, during a usability study for a spreadsheet program, a user interacting with a spreadsheet may be asked to update a graphical chart (e.g., a bar chart) of data on a spreadsheet. The user may have no experience in changing or modifying bar charts once they have been created. The user subsequently exhibits the following behavior: First the user selects the bar chart by placing the mouse pointer on the chart and double clicking the mouse button. Then the user pauses to dwell on the chart for some period of time while introspecting or searching for the next step. Observation of this activity could serve as an indication that the user is experiencing difficulty in using the chart feature of the spreadsheet.

Once the difficulties with using basic functionality are identified, the next step 52 is to build a model for the task of providing intelligent assistance for the basic functionality identified. One type of model is an network of interdependent variables such as a Bayesian network that can represent probabilistic relationships
5 between observable evidence, intermediate variables, and hypotheses of interest. A Bayesian network is a directed acyclic graph in which the nodes represent random variables and arcs between nodes represent probabilistic dependencies among those variables. A Bayesian network represents a joint-probability distribution for a set of variables. Computational procedures, referred to as Bayesian inference algorithms have
10 been developed to operate on Bayesian networks. These algorithms update the probability distributions over states of variables in a Bayesian network given changes that are made in values of states. An example of state changes are new values assigned to states of variables representing observable states of the world, following the observation of these states in the world. A description of Bayesian networks can be
15 found in "Probabilistic Reasoning In Intelligent Systems", by Judea Pearl, Morgan Kaufmann Publishing Inc., San Mateo, 1988.

Some of the nodes are variables that represent evidence that may be observed. Other nodes represent unobservable or difficult to observe hypotheses of interest that may be used in automated decision making. A simple two-variable
20 Bayesian network is displayed in Figure 3. The network represents the dependency between a user being in a state of needing assistance 68 and the user activity of multiple menu exploration 69 or "menu surfing" without making progress. Menu surfing can be sensed, but it cannot be easily observed whether a user needs assistance. In modeling user difficulties with complex software features, evidence nodes may represent user
25 behavior exhibited while performing actions toward certain goals or experiencing difficulty with a software program feature, and hypothesis nodes may represent the goals of the user or the particular assistance needed by the user to complete a task using that feature. For example, while creating a chart the user pauses after selecting the create chart command. An evidence node variable would represent the fact that the user
30 paused after selecting a command. An hypothesis node variable may represent the

probability of the hypothesis that the user needs assistance in creating a chart given the fact that the user paused after selecting the create chart command. Using a Bayesian network is one way of building such a model of uncertain relationships between user actions and hypotheses. Those skilled in the art will recognize that other models can be
5 used such as decision trees or rule-based reasoning.

For any of these models, the next step 54 is to define evidence variables and corresponding modeled events that can discriminate among alternate difficulties and needs for assistance that a person may have in using the basic functionality of the software program, and that can be sensed during the user's interaction with software. A
10 key task required in building the Bayesian network model for the intelligent assistance facility is defining the evidence variables that will be observed, and noting the probabilistic relationships of these variables to variables representing hypotheses of interest. These events are those that are useful in detecting and discriminating among alternate difficulties and needs a person may have in using complex software features in
15 general and, in particular, specific complex software features of a software application.

Classes of observation that are valuable for identifying when and what kind of assistance users need include evidence of search (e.g., exploration of multiple menus), of introspection, (e.g., sudden pause or slowing of command stream), of focus of attention (e.g., selecting and pausing on objects), of undesired effects of action (e.g.,
20 command/undo, dialogue opened and canceled without action), of inefficient command sequences, the structural and semantic content of file being worked on, and goal-specific sequences of actions. Also, program state and program context, user competencies, and the history of information reviewed by the user are useful.

In developing the Bayesian network and specifying these observable
25 variables, it is important to consider the nature and definition of the intelligent user assistance facility which will be harnessed to detect user actions. Consideration must be given to the basic events that can be made available from the software system or application. Modeled events detected as occurring by the system are designed to correspond with each observable variable. The modeled events must be defined in
30 terms of the basic, or atomic events that are generated by the software program.

At run-time, the modeled events corresponding to the observable evidence variables that are modeled in the Bayesian network are provided to the inference system which sets the evidence variables to specific values. The intelligent user assistance facility has the ability to consider the recent occurrence of multiple modeled events in parallel and to consider the time these events occurred, and to use this monitored information to update the state of observed variables in the Bayesian network. The Bayesian network with its associated dynamically changing settings with regard to the status and timing of observed variables is analyzed in the inference system which generates updated likelihoods on different hypotheses of interest. These probabilities are used in making decisions about action and display of information to users.

Upon completing the task of defining evidence variables in a Bayesian network and their associated modeled events, the next step 56 is to identify atomic events that indicate the occurrence of each modeled event. Modeled events can be defined by the occurrence of a single atomic event, associated with a user action or state of the program or may represent a sequence of one or more atomic events. For example, the two atomic events of (1) a user selecting a bar chart and then (2) pausing to dwell on the chart for more than 10 seconds can be composed into a single modeled event that corresponds to an evidence variable in a Bayesian network. Both of these events may be combined to form a modeled event, Sel_Chart_Pause, that indicates an occurrence of a user having difficulty updating a chart with data on a spreadsheet.

Variables in the Bayesian network may be binary and have just two states (e.g., absent and present), or may have multiple states. In cases where there are multiple values, the variable is set to one of the states corresponding to the observation. Binary observational variables are described, but the techniques also apply to variables with greater numbers of values. When the modeled event is observed and is passed to the inference system, the evidence variable corresponding to this event is set to "present" at the time noted for when the event occurred. Whenever probabilistic inference is performed on the Bayesian network, the status of all variables that could be set by modeled events are considered. As a result of inference we may see changes in

hypothesis variables of interest. For example, given the occurrence of Sel_Chart_Pause, and other observations that update the probability that a user is having difficulty with chart-related tasks, the probability assigned to the hypothesis that a user is having difficulty updating the chart with data on a spreadsheet may increase.

5 Another example of a modeled event is the action of a user selecting the Undo command following any one of a set of chart-related commands. The occurrence of one of any atomic or modeled events in a set followed by an undo command would be composed as a Chart_Cmd_Undo modeled event. When this modeled event sets the corresponding evidence node in the Bayesian net, we might see an increase in the
10 probability assigned to the hypothesis that the user may be having difficulty updating chart with data on a spreadsheet. Continuing the example above, atomic events for the modeled event, Sel_Chart_Pause, are (1) a user selects a bar chart, followed by (2) some minimal prespecified period of time (e.g., 10 seconds) in which there is a dwell, indicated a cessation of user activity.

15 Atomic events may serve directly as modeled events. For example, Del_Chart is an atomic event representing the deletion of a chart that is also modeled event, corresponding to the observation variable in the Bayesian network.

Modeled events also include contextual information about the state of data structures in an application. For example, Chart_View, means that a chart is
20 present in the current display. This will increase the probability that charting assistance is relevant to the user.

Words appearing in a free-text query are also atomic events. Higher-level, modeled events may be defined as the appearance of one of a predefined set of words in a user's query. If a user's query is, "How do I change the colors in this chart?,"
25 the modeled event, that one of a set of words (e.g., in this case, a set containing as elements the words "modify," "convert," and "update") has been noted in the query, is passed to the inference system. Other important words in the query, such as "color" and "chart", are also detected as modeled events when they are noted to be in sets of previously modeled sets of similar words. These modeled events are analyzed in the
30 same way as "convert."

The class of modeled events that are persistent across multiple uses of a software program are defined as profile information. These modeled events include events that indicate user competency or lack of competency and are stored in the user profile database. Modeled events stored in the user profile provide additional
5 information about the user, such as the user's previous experience with particular application features or accomplishment of specific tasks, or review of tutorial material in the past. These events may correspond to variables represented in the Bayesian network and used by the inference system to adjust the evaluation of the hypotheses.

Other types of atomic events may include information about the presence
10 or absence of user in front of the computer, as gleaned by proximity detectors, by visual features of the user, including information about the focus of the gaze of the user, expressions of the user, or gestures of the user, as identified by a video camera and vision recognition and tracking system, and speech utterances of a user as detected and processed by a speech understanding system.

After atomic events and their appropriate combination are identified for
15 each of the modeled events, the next step 58 is to construct an event monitor for the software program that will survey user interface actions to extract the atomic events for input to inference analysis. The event monitor constructed is integrated with the software application to instrument that software application at step 60 in making an
20 intelligent user assistance facility. Those skilled in the art will appreciate that the event monitor may be integrated within the application software, or may be a stand alone component that is part of a system library or operating system.

To complete the intelligent user assistance facility, an event processor is
to compose modeled events from atomic user-interface actions (step 62). Next, a
25 knowledge base is constructed from the model (step 64). The last step 66 is to construct an inference engine from the model which uses the knowledge base to determine suggested user assistance for the occurrence of modeled events. A Bayesian network with inference procedures serves both as a knowledge base and inference engine.

Figure 4 is a functional block diagram of the components of an
30 embodiment of an Intelligent User Assistance Facility constructed by using the method

described above. The Intelligent User Assistance Facility includes a software program 72 that has been instrumented to monitor actions or events initiated by a user who is interacting with the software program. Periodically, these actions or events are scrutinized by an event processor 74 to discover whether any of them, either individually or in combination, represent observable variables in the Bayesian network. Modeled events are analyzed by an inference system 76 to form and evaluate multiple hypotheses of what assistance the user may need. These events can lead to increases or decreases in the probability that a user is experiencing a problem in using a particular application feature or in accomplishing a specific task.

The inference system 76 also accesses the user profile system 78 to check for competencies and changes assistance that is given based on user competence. User's background, successful completion of key tasks indicative of competence in particular areas, and previous help reviewed by the user are some of the variables stored in the user profile system 78.

The IUAF controller 79 initiates a cycle of inference analysis by the inference system 76 based upon the setting of its control mode. An inference cycle may be initiated whenever a user asks for help, or whenever one or more of any number of special events have occurred, or may simply occur on a regular clocked cycle, or for any combination of these modes. When the control mode is set to trigger an inference cycle upon occurrence of a special event or the control mode is set to occur on a regular clocked cycle, the inference system 76 will autonomously offer help to the user whenever the computed importance of coming forward exceeds the autonomous assistance threshold set by the user. When the control mode is set to initiate an inference cycle whenever the user asks for help, the inference analysis of user actions and program context is combined with an analysis of the words in a free-text query whenever the user submits a free-text query.

Those skilled in the art will appreciate that the functional blocks illustrated in the diagram may be implemented as separate components or several or all of the functional blocks may be implemented within a single component. For example, a separate library may be implemented that contains the reasoning system functionality

and the event processor functionality. In such an implementation, the reasoning system may query the instrumented program periodically for all of the atomic events since the last query or may simply ask for a specific number of the last occurring events, and then the reasoning system may access the modeled event database for that particular software program to determining whether any of the atomic events are modeled events for input to inference analysis. Those skilled in the art will also appreciate that the inference analysis may be implemented to include only the ongoing analysis of user actions and program context except when a user asks directly for help with a free-text query. In that case, the inference analysis may be implemented to include only the analysis of the words in the free-text query whenever a user asks directly for help with a free-text query. Those skilled in the art will also appreciate that a separate analyses could be undertaken for different classes of modeled events, such as for words in a query and for actions in the interface, and the results of these separate analyses could be combined with the ability to control the weight given to the results of the different analyses.

Instrumented Program

Any software program or library may be an instrumented program, including operating system programs and application programs. As an illustration of an instrumented program, an application program will be described. A typical software application contains a user interface and application procedures. As a user interacts with a software application using input devices such as depressing a key of the keyboard or clicking the button of a mouse, the operating system translates those user actions into events and sends them to the application program. The user interface component of the application program processes those events and forwards them to the appropriate application procedures for responding to those events. The event monitor constructed for the Intelligent user assistance facility for a particular application may be integrated within the application software as illustrated by Figure 5. In this configuration, the event monitor 84 watches for events passing from the user interface 82 component to application procedures 86 and forwards them to the event processor through a call interface to the event processor. It also may be implemented as a stand

alone component that is part of a system or application library which makes periodic calls upon the application program to send it all events since the last period or it may simply request that the application program send only a specific number of the last occurring events. Alternatively, it may be incorporated as part of the operating system in which case the event monitor screens events passing from the operating system to the application program and forwards them to the event processor.

Figure 6 is a functional block diagram of an instrumented spreadsheet application program as in the help-topic embodiment. As a user interacts with a spreadsheet application program using input devices, the operating system translates those user actions into events and sends them to the spreadsheet application program. For example, the user may place the pointer of the mouse on a menu item and may then click on the button of mouse. The operating system will translate that user action into an event that the user has clicked on an application menu item and send the event to the spreadsheet application. The user interface 92 component of the instrumented spreadsheet application program 90 receives that event and forwards it to the appropriate spreadsheet application procedures 96 for processing that event. The spreadsheet events monitor 94, watching for events passing from the user interface 92 component to application procedures 96, will see that event and forward a copy to the event processor.

Event Processor

The event processor is a general monitoring and event composing system that allows high-level events to be created from combinations of low-level events. The event processor does not need to include sophisticated reasoning capability. However, those skilled in the art will also appreciate that a variety of deterministic and probabilistic event preprocessing methods can be employed to work with arbitrary abstractions or clusters of events and to transform events into numeric information with a variety of mathematical functions, including those that might be directly interpreted as likelihoods.

Since the Event Processor produces a set of logical events, any reasoning engine, rule-based or probabilistic, can utilize its output. Figure 7 is a block diagram illustrating the components of the event processor. Atomic events or user input device interactions are sent by an events source 100 such as an instrumented program to the event processor. The instrumented program may use a call interface that specifies the event information to the event processor. Upon receiving these events, the event processor may time-stamp each event and store them in local storage 102, such as system memory, in one of several kinds of data structures for storing information, including a database of records or finite circular queue ordered by the time stamp, or first in, first out order.

A modeled event composer 104 periodically surveys the atomic events in local storage 102 in search of one or more atomic events that may be composed into a modeled event. This typically would be initiated by the inference system 74 at the beginning of an inference cycle; however, the modeled event composer 104 could perform composition of events between inference cycles such as whenever a new atomic event is received, or when special trigger events are received. The modeled event composer 104 then forms modeled events from atomic events according to the event composition rules defined in the event specification. The modeled event composer 104 stores each modeled event that the composer 104 generates from atomic events in the modeled event database 108.

Those skilled in the art will recognize that there could be implemented a modeled event database for each individual program employing the intelligent user assistance facility. For example, there could be a modeled database for each instrumented application program resident on a computer system. Those skilled in the art will also recognize that there may additionally be a separate general modeled database that contains common system or application events such as activating a program or exiting a program. In such a case, the inference system may access both the general modeled event database and the active program's modeled event database during a cycle of inference.

In one example of a modeled event database, there is a single record for every modeled event for that database. When a program with the intelligent user interface is activated, pre-existing persistent events stored in the user profile may be used to initialize the modeled event database. As modeled events are formed by the composer, the record for that event is updated in the modeled event database. An example of a record format for modeled events in this modeled event database is:

Fields Name	Type	Description
EventName	string	Name of modeled event
Occurred	boolean	Flag which indicates that the modeled event occurred within its horizon.
Atomic	boolean	Flag which indicates that the modeled event is derived from a single atomic event.
StartTime	real	Lowest start time on any atomic events from which this modeled event is composed.
EndTime	real	Highest end time on any atomic events from which this modeled event is composed.
Value	real	This field holds the result of the evaluation of an expression
LastTime	real	The time of last occurrence of the modeled event
Count	real	Total number of occurrences over all sessions
LastTimeThisSession	real	Time of last occurrence this session
CountThisSession	real	Number of occurrences this session
Persistent	boolean	Flag which indicates that this is a persistent modeled event.

Of the fields used to count the number of occurrences and the rate of their occurrence, LastTime and Count are stored in the user profile database for all persistent modeled events. Upon initialization of the program, these fields get loaded from the user profile database into their corresponding modeled events in the modeled event database.

Modeled events are only relevant for a comparatively short time. This period of time during which a modeled event is relevant may be referred to as its horizon. Different events may have different horizons, a sixty second horizon is a suitable length of time for the horizon of most events generated by user actions. The duration of the horizon determines the length of time that a modeled event is considered in the inference system 76.

Each modeled event may be assigned a default system or application horizon or it may be assigned a specific horizon. Modeled events that have resided in the modeled event database 104 beyond their horizon time are discarded whenever the modeled event database 104 is accessed or updated. However, before they are discarded from the modeled event database 104, persistent modeled events are recorded in the user profile. The inference engine 76 and user profiler, in general, act as event sinks and access the modeled event database 104 to extract modeled event records. However, there is an exception. At program initialization time, pre-existing persistent events stored in the user profile may be used to initialize the modeled event database 104 and, as a consequence, modify the program's behavior.

Event System Specification Tool

The event specification tool allows for rapid creation of the modeled event composer and modeled event database for specific application programs during creation of an intelligent user assistance facility for that software program. A language was developed for constructing modeled events (e.g., menu surfing) from atomic events made available by a program, for instance an application (e.g., multiple menus visited without action within 10 seconds, position of cursor on spreadsheet, etc.). The language allows atomic events to be used as modeled events directly, as well as for atomic events

to be formed into higher-level modeled events. Also, modeled events can be further combined to create other modeled events. More specifically, the language allows for Boolean combinations of atomic events and for combining the Boolean events into modeled events with operators that capture notions of temporality. The event specification tool can be linked directly with a Bayesian network modeling tool so that the modeled events and their definitions can be generated while building and refining the Bayesian user model.

Figure 7 is a block diagram of an event system specification tool 110 used for generating code for the modeled event composer 104 and modeled event database definitions 118. Each modeled event definition 112 created during the step of defining modeled events is used to specify the events in the event language 114. The event specification 114 is then input to an event language translator 116 which translates the event language statements of the specification into high-level language code, such as C, to create the modeled event composer 104 component. The event language interpreter also creates the modeled event database definition. As part of that process, the event language interpreter 116 may also create the modeled event database 106 directly from the modeled event database definitions 118 in systems employing a database structure to store information on the status of modeled events.

The event language includes Boolean and set operators, with the addition of temporal operations including dwell and sequence. Modeled events are defined in a separate text module written in the event specification language. However, they can also be constructed with the use of special definition forms. The text definitions are created as named statement blocks for each modeled event. These definitions contain the labels of atomic events and indicate operations on one or more atomic events.

The general form of a modeled event declaration in the event definition language is:

```
declare ModeledEventName
{
  <expression list>
}
```

An example of the modeled event declaration for the Dwell_on_Chart modeled event is:

```
5      declare Dwell_on_Chart
      {
      (and (Select_Chart, Dwell(5,s))
      }
```

The interpreter translates the syntax statement of the event definition language to the modeled event name, Dwell_on_Chart and updates the modeled event database by adding a record with that event name and initializes all record fields. A modeled event may be a complex synthesis of several other events, including other modeled events, or it may be simply a renaming of an atomic event. The event synthesis language supports infix expressions which entail logical combinations of atomic events, modeled events, arithmetic, and logical operators. As each elementary operation is performed, the intermediate result is treated as an event and has similar properties.

The event language has logical, set-theoretic operators, temporal, and arithmetic operators. Operators allow the definition of specific events as well as events defined as generalizations based on sets of events. For example, higher-level events can be built via abstraction or generalization with operators like Element:

25 Element($\{x_1, \dots, x_n\}$): Any event drawn from the elements of the set of events $\{x_1, \dots, x_n\}$ occurs. Such sets can be defined as specific classes of events.

Temporal operators allow us to define events in terms of sequences of events over time. Temporal intervals may be measured in terms of the number of user actions, the amount of time, or other measures of duration that can be defined as functions of user actions and time. For example, it can be useful to define modeled events in terms of Scaled_seconds, a measure of duration where seconds are scaled by the rapidity at which a user is working, so as to adapt the temporal constants to users that work at faster or slower rates than an average user. When specifying temporal operators, we indicate the dimension of the measure of temporal interval by specifying

the temporal dimension d with a value (e.g., c for commands, s for seconds, and ss for scaled seconds). Some useful temporal operators are:

Rate(x, t, d): At least x atomic events occur in duration t of dimension d .

5

Ensemble_N($\{x_1, \dots, x_n\}, t, d$) At least N events from the set of events $\{x_1, \dots, x_n\}$ occurs within interval t of dimension d .

All($\{x_1, \dots, x_n\}, t, d$): All events in a specified set of events $\{x_1, \dots, x_n\}$ occur within duration t of dimension d .

10

Sequence($\{x_1, \dots, x_n\}, t, d$): All of the events specified in a set of events $\{x_1, \dots, x_n\}$ occur within duration t of dimension d .

TightSeq($\{x_1, \dots, x_n\}, t$): All of the events specified in a set of events $\{x_1, \dots, x_n\}$ occur within occur within duration t of dimension d .

15

Dwell(t, d): No user action for t seconds or scaled seconds of dimension d .

In addition, parameters can be specified that define the persistence or dynamics of probabilistic relationships between modeled events and other variables in a Bayesian network with increasing amounts of time since the modeled event has occurred. In the general case, we can provide any temporal function that describes how the probabilistic relationships of events to other variables in a Bayesian network change as the time following the occurrence of an event increases.

20

More specifically, the language allows the specification of how $p(E_i, t_p | H_j, t_p)$, the probability of event E_i occurring at time t_p in the past conditioned on the truth of hypotheses H_j at the present moment, t_p changes with the increasing amount of time $t_p = t$ since the event last occurred.

25

Useful distinctions for representing the temporal dynamics include horizon and decay:

Horizon: The interval of time beginning at the time $t=t_0$ that an event becomes true, that the probabilistic relationship persists as $p(E_i, t_0 | H_j, t_0)$ without change. The probability remains unchanged as long as difference between the present moment, t_p , and the time, t_0 , the event was most recently seen is less than the horizon.

30

Dynamic: The time-dependent functions that describe $p(E_i, t_o | H_j, t_p)$ and $p(E_i, t_o | \text{not}(H_j), t_p)$ after a horizon is reached. The dynamics can be summarized with a description of the change in the likelihood ratio defined by the ratio of these probabilities. The functions indicate how the initial probabilities $p(E_i, t_o | H_j, t_p)$ and $p(E_i, t_o | \text{not}(H_j), t_p)$ change as the time the event became true becomes increasingly distant from the present time, t_p . It can be asserted that these probabilities converge at some time to $p(\text{not}(E_i) | H_j, t_p)$ and $p(\text{not}(E_i) | \text{not}(H_j), t_p)$, and the initial likelihood ratio, when $t_p = t_o$, converges to the ratio of these latter probabilities, and, therefore, functions can be assessed that decay the probabilities or likelihood ratio from the initial likelihood ratio when $t_p = t_o$ to the likelihood ratio associated with absence of the event.

The probabilities can be computed as a function that takes as arguments the initial probabilities (at t_0), the probabilities associated with the finding being absent, and the time of the present moment, t_p , and ensures consistency among probabilities at any moment. That is,

$$p(E_i, t_o | H_j, t_p) = f [p(E_i, t_o | H_j, t_o), p(\text{not}(E_i) | H_j, t_p), t_p]$$

and

$$p(E_i, t_o | \text{not}(H_j), t_p) = f [p(E_i, t_o | \text{not}(H_j), t_o), p(\text{not}(E_i) | \text{not}(H_j), t_p), t_p]$$

Operators being used to define modeled events may affect the duration information of the resulting modeled event. Conjunction results in storing the latest occurrence time of any subordinate events. Disjunction results in storing the times associated with the first true subordinate event.

The Value field of the generated intermediate result of the modeled event is set according to the result of the expression. The result of all arithmetic comparison operations is either one (Boolean TRUE) or zero (Boolean FALSE). If the argument of an arithmetic operation is an event, the argument is evaluated as one if the event has occurred and zero if the event has not occurred. The time stamps associated with simple numbers are set to the span of the current cycle; in other words, the start time is the time of the last cycle and the end time is the current time.

Since the syntax allows for recursion, care should be taken that infinite recursion not occur. These statements are entered into a text file known as the event specification module.

The event specification module 114 is then processed by the event definition language translator 116. The event language translator converts the event specification language 114 into C++ code which may be compiled to machine language and linked directly to the event processor system library, IUAF library, operating system library, or other implementation library. The results of translating a high-level event specification is the modeled event composer 104 which creates modeled events.

10 When the modeled event composer 104 is invoked during a cycle of inference, it builds the set of currently active modeled events and then returns to the caller. The inference code then performs appropriate actions indicated by the event set.

User Profile System

15 The inference system 76 accesses the user profile system 78 to check for competencies and changes assistance that is given based on user competence. User's background, successful completion of key, tasks indicative of competence in particular areas, and previous help reviewed by the user are variables that can be stored in a persistent file and updated with time. Such persistent information about user

20 background, experience, and competence is referred to as "profile information."

Figure 9 is a block diagram of the user profile system 78. The user profile system 78 includes a user profile database 120, a user profile access routine 122, an event status update routine 124, a background and competency dialog routine 126, and a customized tutorial routine 128.

25 The user profile database 120 contains records identical in format to the records in the modeled event database 106 so that there is a common format between the modeled event database records and the user profile database records. This common format makes it more convenient both for updating the user profile database with modeled events and for the inference engine to interpret these records when it access

30 these databases for inference analysis. The user profile database records are stored

persistently so that they create an historical record of particular user's competency with specific software applications. This database is maintained on a by-user-by-application basis. An example of persistently stored modeled events that are indicative of user competency are the completion or non-completion of certain tasks, successful or unsuccessful use of particular features, and assistance received or help information reviewed in the past.

Figure 10 is a flowchart illustrating how the user profile database 120 is updated with event status of a user interacting with a specific application program. The event status update routine 124 in step 130 retrieves the first modeled event record in the modeled event database and checks in step 132 whether the record is marked as a persistent record. If it is marked as a persistent record, then the record is added to the user profile database in step 134. If it is not marked as a persistent record, then it is not added to the user profile database. Each subsequent record in the modeled event record database is retrieved in turn (steps 136 and 138), checked whether the record is marked as a persistent record (step 132), and added to the user profile database if it is marked as a persistent record (step 134).

The availability of a standard stored user profile containing information about "persistent events" could follow a user around over a local network or over the Internet, custom-tailoring the user's software wherever the user may be. While a software application is active, a new user may log on to the system and use the active software application. Figure 11 is a flowchart illustrating how a user profile is accessed whenever a new user signs on to the system or application. When a user logs on to the system or application (step 140), the user profile access routine 122 searches the local user profile database for that user's profile (step 141). If the profile is found, the path to that user profile is passed to the inference system 76 (step 142). If that user's profile is not found, because that user may be using the application at a remote location, then the user profile access routine searches the available networks for any user profile databases that contain the user profile (step 143). If the user's profile databases is not found or the user's profile is not found in any user profile databases accessed, then the user profile access routine asks the user for permission to enter a user competency dialog with the

user (step 144) and creates a user profile for that user (step 145). The path for that user profile is passed to the inference system 76 (step 142). If the dialog is not completed, the path to default profile information is passed to the inference system 76 (step 146) and used.

5 In addition to maintaining a persistent user profile for use during real-time inference, a continuing background analysis of repetitive patterns of a user's needs for assistance during one or more sessions that exceed a threshold, may be used to customize a tutorial of help information that may be offered to the user at the end of a user's session with the that program. Figure 22 is a flowchart of the routine for
10 generating a customized tutorial. During a user's session with a specific application program, a histogram of help topics exceeding a relevant probability threshold is recorded by the inference system 76 (step 150). When the user asks to quit the session with the application program, the user profile saves the histogram in the user's profile (step 152) and prepares a customized tutorial of help information related to the help
15 topics recorded in the histogram (step 154). The user profile system notifies the user upon exiting the application program that a customized tutorial was prepared for online review or for printing as a manual (step 156).

Inference System

20 When a cycle of inference analysis is initiated, the inference system 76 accesses the modeled event database 106 to extract modeled event records for each of the modeled events that has occurred since the last cycle of inference. A functional block diagram of the inference system is displayed in Figure 13. The inference system 76 includes one or more knowledge bases 160 and an inference engine 165. Knowledge
25 bases 160 include information relating variables that represent observable states of the world, such as user actions and words, to variables that represent hypotheses of interest about a user's goals and needs for assistance that may not be observed directly or that are costly to observe. For example, the goals of users can not necessarily be directly inspected, but a sequence of computer commands can be recorded. A user can be
30 actively queried about goals, but this may be quite distracting to users.

A deterministic knowledge base consisting of an interrelated set of logical rules may link observations to hidden hypotheses. However, it is often more appropriate to process the uncertain relationships as probabilities between observations and the likelihood of hypotheses about a user's needs. For example, in attempting to understand and to predict the behavior of a complex system such as human physiology it is typically not possible to completely model with deterministic relationships all components of the system, and to then have access to a deterministic model for performing diagnosis or designing therapy based on a set of symptoms. For diagnosis and decision making about complex systems, we are often forced to reason under uncertainty, and to explicitly address the incompleteness in our understanding. Probability provides us with a means of diagnosis and forecasting about the behavior of complex systems given knowledge about a set of relationships among observational variables and hidden variables we may identify as being important. We can use probabilistic methods to represent and reason about weak and strong uncertain dependencies among observations, such as symptoms of a patient and variables such as diseases in a patient. A good example of complex system is a user attempting to perform a task while interacting with a computer software application or system. It is very difficult to build deterministic models that link a user's behavior to a user's goals and intentions. The best way to diagnosis a user's needs is to develop appropriate abstractions, based on our understanding of the relationships among various kinds of user's goals and actions, and to represent and reason about the relationships with probability. Probabilistic methods allow us to build models at a level of abstraction that is appropriate in light of our incomplete understanding about users and their actions.

A Bayesian network or a generalization of a Bayesian network, called an influence diagram, can be employed to represent the certain or uncertain relationships among user actions and such hidden, but important, states as user goals, user intentions, and user needs, given observable information such as one or more user actions. A Bayesian network is a directed acyclic graph where nodes are random variables and arcs represent probabilistic dependencies among those variables. Variables in a Bayesian network are chance variables or deterministic variables. A Bayesian network represents

a joint-probability distribution for the set of variables it represents. The probability distributions over the values of chance variables depend on the values of direct ancestors, or parent variables that are parents of nodes. The value of deterministic variables is a deterministic function of predecessors. Influence diagrams are a
5 generalization of Bayesian networks which represent additional nodes that represent possible actions and the utility of outcomes.

The inference engine 165 includes inference procedures 166 that operate on the knowledge base 160. The knowledge base 160 includes a Bayesian network 162. Those skilled in the art will recognize that the knowledge base may be a deterministic
10 knowledge base with logical chaining procedures as the inference procedures of the inference engine. Or the knowledge base may be a Bayesian influence diagrams with inference procedures for operating on Bayesian networks as the inference procedures of the inference engine. Temporal reasoning procedures 167 and value-of-information procedures 168 are also included as part of the inference engine 165. Other specialized
15 procedures may be included.

Bayesian inference procedures operate on Bayesian networks to compute a consistent posterior probability distribution over the values of variables in the network given the setting of states of observable, or evidence, variables to specific states based on the observations. Assume we are interested in the probability distribution over
20 unobserved hypotheses of interest, H_1, \dots, H_m , and have access to observations and profile information E_1, \dots, E_n . Bayesian network inference algorithms compute a probability distribution over H given the observations, written $p(H | E_1, \dots, E_n)$. For influence diagrams, decision-theoretic inference is performed to identify the expected utility of alternative actions. The best action is the one associated with the highest expected
25 utility.

In the most general sense, the task of designing inference for an intelligent user assistance facility is best undertaken from the perspective provided by decision theory, an extension of probability theory to reflect concerns of value under uncertainty. Influence diagrams allow us to express fundamental relationships about
30 uncertainty, action, outcomes following action, and the value of those outcomes. This

representation is useful in the design, understanding, and in many cases, the actual implementations of intelligent user assistance systems. However, it is often the case, that a simpler systems may be built without explicitly representing decisions, outcomes, and the utility of outcomes, and instead to use Bayesian networks to represent probabilistic relationships, and to use probabilistic inference in combination with procedural controls and thresholds as approximations of more complex influence diagram models.

A general influence diagram for user modeling and action to assist users of software is portrayed in Figure 14 as an exemplary embodiment of a Bayesian influence diagram as a knowledge base. As portrayed in the figure, a user's background 170 influences with uncertainty, a user's overall goals 171 in using software, as well as the user's knowledge 172 in using software. User's knowledge 172 is also influenced by the previous help 173 that user may have seen. The user's background 170 and previous help 173 are variables that can be stored in a persistent file and updated with time. Such persistent information about user background, experience, and competence is referred to as "profile information." As indicated in the influence diagram, the user's goals 171 and knowledge 172 in turn influence with uncertainty, the informational needs 174 of the user. The goals 171 and needs 174 in turn influence the sensed activity 175 and the words that might be used in a query 176 to the software, or software's help system. A utility function of the user 177 is represented as a diamond. The utility 177 is influenced directly by the informational needs 174 of the user, the cost of taking autonomous action 178 (e.g., distraction for the user's current focus of attention) and the action that is taken. Several classes of action may be available to the system, including the providing of advice or help 179, the execution of software actions of various kinds 180 and the acquisition of additional, previously unobserved information from the system, 181 or directly from the user. As indicated in the Figure 14 by the scale on the cost of assistance variable 178, we may wish to allow the user to directly change the cost of the assistance so as to control, in a fluid manner, the degree of autonomy given to the system.

The overall goal of the inference system is to identify actions that will optimize the user's expected utility given the user's needs and the cost of taking autonomous action. Given a set of evidence E about a user's background and actions, the probability distribution over hypotheses of interest about a user's needs, H, is computed. In the case of the Bayesian influence diagram in Figure 14, the probability distribution over a user's needs for assistance is computed. The expected value of actions in the set of all possible actions A must be considered. To do this, consider the utilities over outcomes, represented as a utility model (diamond node). The utility model contains information about the value or utility of outcomes. Outcomes (A,H) are defined as doublets of actions A taken and the state actual user's needs H. The utility model tells us the utility associated with each outcome, $u(A,H)$. The best action to take, A^* , is the action that maximizes the expected utility under uncertainty, computed as follows:

$$A^* = \arg \max_A \sum_j u(A, H_j) p(H_j|E)$$

Although influence diagrams represent explicit models of action, Bayesian networks are often easier to build and can be used for decision making and action by employing thresholds or rules about probability over variables in the network to indicate when actions should be taken based on the degree of belief assigned to a variety of states including states that describe preference about action. Influence diagrams such as the one portrayed in Figure 14 can be used to clarify approximations used in the Bayesian network.

25 Temporal Reasoning and Dynamics

The inference system 76 also contains special temporal reasoning knowledge 164 in the knowledge base 160, and procedures 167 for performing inference about the changing relationships of observations to other variables in the model as the observations occur at progressively more distant times in the past. In the general case, applications of Bayesian procedures and knowledge bases to reasoning about patterns of observations seen over time requires the consideration of variables and

their interdependencies within a single time and at different times. Building and performing inference with Bayesian networks that include copies of variables and dependencies for different times slides can lead to difficult computational reasoning problems and unacceptable response times for a user interacting with a program. Thus, several approximations are of value. In one kind of approximation, only specific dependencies considered to be of the greatest importance are considered over time.

A portion of a Bayesian network with sets of variables at different times and prototypical dependencies represented is displayed in Figure 15. Figure 15 displays a Bayesian network that explicitly represents random variables at different points in time, and probabilistic dependencies among variables at a single time and among variables at different times. The lack of dependencies among variables are assertions of assumptions of independence. The figure displays possible influences among the primary relevant assistance to provide a user at present moment with the most relevant assistance to provide users at earlier and later times. Also, the figure highlights potential dependencies among the status of observations (E_i and E_j) in the present, with observations made in the past and future. The figure also indicates the relationship between the primary relevant assistance in the present and observations in the past. In the general case, such multiple-connected networks are difficult to solve and to assess during the construction of the models.

We will now describe an approach to reasoning about relevant assistance over time that is more tractable to solve and to assess, through making additional assumptions of independence and representing the relationships among evidence seen at progressively more distant times in the past and hypotheses of interest with parameterized functions that dictate the strength of the probabilistic relationship between those observations and current the goals or needs for assistance. The use of this approximation in an exemplary embodiment employs a single explicit Bayesian network, but allows the system to consider implicitly multiple Bayesian networks over time. This method takes advantage of direct assessment during modeling time of functions that adequately describe the dynamics of the probabilistic relationships between observed variables and other variables in a Bayesian network. At modeling

time, a horizon and dynamics are assessed for each observational variable E_i , and the variable is annotated with this information. At run-time, the inference engine 165 includes procedures 167 for using this information about horizon and dynamic. The horizon for an event captures the interval of time the probabilistic relationships with hypothesis H_j persists as $p(E_i, t_0 | H_j, t_0)$ and $p(E_i, t_0 | \text{not}(H_j), t_0)$ without change. The dynamics is the time-dependent change of the probabilities in the present moment as described earlier.

Heuristics in defining a horizon for events include the use of a default horizon for all observational events considered by a system unless otherwise specified, and the use of an event queue of finite length k , where only the last k modeled events are considered in the analysis. This approach can be combined with dynamics as just described.

Value of Information

The inference engine 165 also includes expected value of information (EVI) procedures 168 to compute the expected value of acquiring information from a user. EVI procedures and information-theoretic approximations that can be employed in probabilistic models are well-known to people skilled in the art of decision theory. The EVI is a means for computing the expected value of acquiring information about variables that have not yet been observed. Such information includes answers by users to questions posed by the computer system about their goals and needs. The Net Value of information (NEVI) is net value of gathering information, including the costs of gathering the information. The inference system 76 only asks a question, or gathers information autonomously, when the informational benefits of the information outweigh the costs. If E_x refers to each previously unobserved variable that can be queried, O_k is a possible observational value of E_x , and A is the value that will be observed when the variable is evaluated, the NEVI of each variable is:

$$\begin{aligned} \text{NEVI}(E_x) = & \sum_k p(E_x = O_k | E) * [\max_A \sum_j u(A_i, H_j) p(H_j | E, E_x = O_k)] \\ & - \max_A \sum_j u(A_i, H_j) p(H_j | E) - C(E_x) \end{aligned}$$

where $C(E_x)$ is the cost of evaluating evidence E_x . EVI can be applied directly to influence diagrams. Several well-known information-theoretic algorithms, including those based on computation of entropy, function to provide similar functionality for decisions about information gathering in probabilistic models such as Bayesian networks. These algorithms are typically more tractable than computing NEVI with influence diagrams. They do not explicitly manipulate utilities, but nevertheless can be employed in conjunction with heuristics that balance the cost of information with measures of informational value.

Figure 16 is a flowchart of value-of-information procedures in the value of information component. The value of information component lists all of the previous unobserved information in order of the value of it as alternative information (step 190) and weighs the benefits of obtaining the information against the cost of distracting the user for each item of information (step 192). If the benefit is greater than the cost set by the user, the IUAF asks the user to supply the item of information that has the highest value (step 194). If the user responds, then the value of information component repeats the first two steps until either the cost of distraction is greater than the benefits of obtaining the information or the user does not respond (step 196).

Scenarios Illustrating the Screen Outputs of the IUAF

Figures 17 through 23 illustrate the screen outputs of the intelligent user assistance facility in an exemplary embodiment of a spreadsheet application. The system considers a user profile, user-specified threshold information, and combines words and user actions to provide intelligent assistance to the user. The system has the ability to respond to requests for assistance as well as to autonomously provide assistance. Figure 17 shows an assistance interface 200 that is generated when the user explicitly has requested help. Guesses about useful assistance to provide the user are displayed in the list box 201 in the upper left-hand corner. The current system settings box 202 in the lower left-hand corner indicates the user is an expert user. This information was provided by the user profile system. The current system setting box 202 also indicates that both user actions and free-text queries are used by the inference

engine during a cycle of inference analysis. Given the profile information and the recent actions involving interaction with the layout of text, in combination with pauses, and with the user taking action to change the appearance of rows and columns shortly before help was actively requested, the system believes the user may best benefit by receiving assistance on working with fonts, advanced formatting instructions, and with methods for changing the dimensions of the spreadsheet document. To the right of this screen photo is displayed an inference graph 203 showing the probability distribution over a set of task areas that was generated by the Inference System 76. This display was built for engineering purposes but may be displayed to the user. The length of the bar graph next to each task area is the probability assigned to that area. The two highest probability areas (those with the longest associated bar charts) are for the tasks of working with fonts and advanced formatting topics.

In Figure 18, the profile information was changed to a profile for a novice user. The expertise level in the current systems setting box 202 indicates that the user has a level of expertise of a novice. All other events handling remained unchanged. We have used a profile for a novice user. The profile information is used to modify the probabilistic relationships in the Bayesian network. The inference graph 203 displaying the probability distribution over relevant assistance shows the revised probability distribution with the profile information. As can be seen in the best guesses list box 201, the system now believes that the user may most benefit by assistance with basic working with rows and columns of the spreadsheet, changing alignment and basic formatting of charts.

Figure 19 shows the analysis of the same set of user actions, but now, with consideration of the words in a user's query. The user has input to the system the natural-language query, "How do I make this look prettier?" in the query input box 204. The analysis of the words and the actions are combined to generate a new probability distribution, as indicated in the inference display 203, and a new corresponding list of recommended topics in the best guesses list box 201. Now, the list is resorted and updated with autoformatting assistance being recommended as the most relevant

assistance, but also still contains other formatting topics, including changing alignment, working with borders, and working with fonts.

Figure 20 shows the program at a later time when help is again requested after interaction with a chart in the document, coupled with pauses on chart commands.

5 The probability distribution computed by the inference system 76 is displayed in the inference graph 203. Now, the recommended topics in the best guesses list box 201 focus on a variety of high-likelihood areas in the area of graphing information. Figure 21 shows the addition of information submitted in the form of a query in the query input box 204. The natural language query from the user is, "I need to access data from a
10 different application". The words are now considered in conjunction with the events. The inference system 76 provides a revised probability distribution in the inference graph 203 and a revised list of recommended topics in the best guesses list box 201, centering in retrieving data from a database and about changing the data displayed in a chart.

15 Figure 22 illustrates the screen outputs of the inference system 76 offering autonomous help to a user. While the user has been interacting with a spreadsheet program 205, an assistance monitoring agent computes the probability that a user needs assistance and the inference system 76 also computes the likelihood of alternative kinds of help to give the user, should the user require assistance. The
20 assistance monitoring agent displays the result of inference on the probability that a user needs assistance in the assistance monitoring agent window 206. Like the inference graph 203 display described above, the assistance monitoring agent 206 was built for engineering purposes but may be displayed to the user. The user who has been interacting with a spreadsheet program 205, now selects the whole sheet and pauses. As
25 the assistance monitoring agent window 206 shows, the computed probability that the user needs help moves from 26% to 89%, reaching a threshold, and a small timed window, the autonomous assistance window 207, autonomously comes forward displaying the highest likelihood assistance areas that were computed. In this case, the highest likelihood topics include formatting cells, checking for spelling errors, and
30 performing calculations, topics that have implications for analysis or modification of

the entire document. The autonomous assistance window 207 politely offers help. It also provides a chance for the user to reset a threshold that will increase or decrease the probability required on the likelihood the user needs help before displaying the window. The window 207 will time out and be removed with an apology for distracting the user
5 if the user does not interact with the window 207.

Figure 23 illustrates the display 208 of the topics that have been determined to be of value to the user based on the recording and summarization of ongoing background inference about problems that user has been having during a session using a software program. Profile information about the user's background,
10 competence, and about the assistance that the user has received or reviewed in the past are also included. These topics have been prepared for printing as a custom-tailored manual for later perusal.

Bayesian Network Annotated with Temporal Reasoning

Figure 24 shows an example of a Bayesian network that computes the probabilities of alternative forms of assistance to provide a user as a function of profile information about the user, as well as recent actions taken by the user. To indicate that this structure may be repeated for many classes of observations and task-specific needs for assistance, variables x, y, z, and A have been used in several of the nodes.
15

In this model, we represent profile variables that represent the status of user competency in two different areas of software functionality, x and y, with a set of variables labeled User Competence x 210 and User Competence y 211. The user competencies are each influenced by variables representing (1) specific observed activities by the user, (2) the history of help topics that have been reviewed by the user,
20 and (3) the overall user background 212 that may be uncertain or set by a dialogue with the user. As indicated in the Bayesian network, the user competencies directly influence the prior probability of different assistance being relevant to the user, as represented by the states of the Primary Assistance Needed variable 213. This variable is also influenced by a context variable 214, that describes information about program
25 state, such as the existence in the current version of an application of a particular data
30

structure. In this model, the problem is formulated with a single primary assistance variable, with states representing alternate forms of assistance. The model assumes that there is only one state of primary assistance the user needs at any particular time. Multiple variables about assistance can also be represented to capture the notion that
5 several types of assistance may all be relevant at a single time.

As depicted in the Bayesian network, the Primary Assistance Needed node 213 influences a set of observed modeled events, such as selecting a graphical object Z, then pausing for more than some predefined period of time (Sel Z Pause 215), or modifying an object and then performing an undo, (Z Cmd Undo 216). At run time,
10 these modeled events are detected and passed to the Inference System 76 by the Event Processor 74. The corresponding observable variables to the modeled events are update and Bayesian inference is performed to update the probabilities over all of the unobserved variables in the system, including the Primary Assistance Needed variable 213. The time of the event is noted and as the event flows progressively into the past,
15 temporal reasoning procedures 167 are applied to update the probabilistic relationships in the network.

Figure 25 shows an instantiation of the nodes with variables x and z with distinctions with relevance to reasoning about the relationship of actions and profile information on charting or graphing information from a spreadsheet, such as the
20 Microsoft Excel product. Variables that are persistent and stored in a profile are marked with an adjacent P. Variables that are observed as modeled observable events are marked with M. One or more nodes may represent various classes of competency with charting. In this case, a single variable, User Competence Charting 220, is displayed as the variable that contains information about the overall level of charting ability of the
25 user. The probability distribution over the states of this variable are influenced by Hx Help Charting 221 and Hx Activity Charting 222, and User Background 212. Hx Help Charting 221 is set by observations that are stored in persistent profile about one or more charting activities the user has completed successfully and information about the help information on charting in an online manual that the user has reviewed.
30 Competencies about charting and a variety of other areas of software functionality

influence the Primary Assistance Needed variable 213. This variable in turn influences the probability of observing different behaviors by the user. For example, the state of a user needing assistance on a specific type of charting task, influences the probability of seeing the user select and pause for s seconds on a chart in the document (Sel Chart
 5 Pause 223). Appendix I contains a more complete Bayesian network in an exemplary embodiment of a spreadsheet application.

Figure 26 represents an exemplary embodiment of temporal knowledge and inference. Each observable modeled event is annotated with information about the effects of time flow on the probabilistic relationships of the event with other variables in
 10 the Bayesian network. In this case, each node is annotated with temporal dynamics information about a horizon and decay, which may be implemented as a table of probabilities or other data structure. The dimensions of time for the temporal dynamics of each variable are indicated in seconds that have occurred since the event became true.

We now provide an example of a case based on components in the
 15 knowledge base displayed in Figure 26.

Let us consider a case where we have the following profile information:

User_Background 212: Consumer_Novice
 HX_Help_Charting 221: No_Help_Reviewed
 20 Context: Normal_View
 Objects:Sheet(1),Chart(2)

At time t=50 seconds, we have noted the following modeled events:

25 Chart_Cmd_Undo 224: Absent
 Chart_Create_Pause 225: Absent
 Sel_Chart_Pause 223: Absent
 Move_Chart_Pause 226: Absent
 Chart_Dialog_UnSucc 227: Absent
 30 Menu_Surfing 228: Absent
 Click_Unrelated 229: Absent

35 Sel_Chart_Pause 223

$$p(E_1, t_0 | H_j, t_0) = .15$$

$$p(\text{not}(E_1) | H_j, t_0) = .85$$

$$5 \quad p(E_1, t_0 | \text{not}(H_j), t_0) = .005$$

$$p(\text{not}(E_1), t_0 | \text{not}(H_j), t_0) = .995$$

10 Horizon: 5 seconds

Dynamics:

$p(E_1, t_0 | H_j, t_0)$: linear convergence at 15 seconds

15 $p(E_1, t_0 | \text{not}(H_j), t_0)$: linear convergence at 15 seconds

20 Chart_Dialog_Unsucc 227

$$p(E_2, t_0 | H_j, t_0) = .08$$

$$p(\text{not}(E_2) | H_j, t_0) = .92$$

25

$$p(E_2, t_0 | \text{not}(H_j), t_0) = .0005$$

$$p(\text{not}(E_2), t_0 | \text{not}(H_j), t_0) = .9995$$

30 Horizon 20 seconds

Dynamics: 0

35

A cycle of analysis is initiated at $t=50$ seconds. The cycle includes the composition of modeled events from atomic events local storage 102, the transmitting of the events to the Inference System 76, and inference with the new observations. We will explore the current probability that the primary assistance needed by the user is help with charting information, $p(\text{Charting_Assistance}|E)$. The prior probability of the user needed help in this area conditioned on the user background, competence, and history of reviewing help is .001. The context information has increased this to a

40

probability of .04. Probabilities are also assigned to other hypotheses about the primary assistance as a function of all of the evidence seen so far (but that is not explicit in this example).

At t=52 seconds the user selects a chart. The atomic event of chart being
 5 selected is detected and is stored in the atomic event local storage 102. The user then pauses for 3 seconds without doing anything. A dwell of 3 seconds is detected in the event local storage 102. A cycle of analysis is called at t=60 seconds. A modeled event Sel_Chart_Pause 223 is composed by the Event Processor 74, tagged as becoming true at t=55 seconds, and is sent to the Inference System 76.

10 The Inference System accesses the Bayesian network that is annotated with temporal dynamics information. As indicated in the table of probabilities for Sel_Chart_Pause 230, there is a horizon of 5 seconds, and a linear convergence of probabilities to the probabilities of absent at 15 seconds.

At t=63 seconds, the user selects a chart dialog and cancels the dialog
 15 with out success at t=66 seconds. A cycle of analysis is called at 70 seconds, and the modeled event Chart_Dialog_Unsucc 227 is composed and passed to the Inference System 76 with a time stamp of 66 seconds. As indicated in the table of probabilities for Chart_Dialog_Unsucc 231, there is a horizon of 20 seconds, and no temporal dynamics.

20 This user has set the threshold for being bothered with automated assistance at $p=.12$. Thus, the user is offered assistance on charting when the probability that this topic is relevant rises to $p=.14$.

25

At t = 50 seconds

30 $p(\text{Sel_Chart_Pause}, t_0 | H_j, t_p) = .15$

$p(\text{Sel_Chart_Pause}, t_0 | \text{not}(H_j), t_p) = .005$

$p(\text{not}(\text{Sel_Chart_Pause}) | H_j, t_p) = .85$

$$p(\text{not}(\text{Sel_Chart_Pause})|\text{not}(H_i), t_p) = .995$$

5

$$p(\text{Chart_Dialog_Unsucc}, t_o | H_i, t_p) = .08$$

10

$$p(\text{Chart_Dialog_Unsucc}, t_o | \text{not}(H_i), t_p) = .0005$$

$$p(\text{not}(\text{Chart_Dialog_Unsucc})|\text{not}(H_i), t_p) = .9995$$

$$p(\text{not}(\text{Chart_Dialog_Unsucc})| H_i, t_p) = .92$$

15

Inference result: $p(\text{Charting_Assistance}|E, t_p) = .02$

20

At $t = 60$ seconds

$$p(\text{Sel_Chart_Pause}, t_o | H_i, t_p) = .15$$

25

$$p(\text{Sel_Chart_Pause}, t_o | \text{not}(H_i), t_p) = .005$$

$$p(\text{not}(\text{Sel_Chart_Pause})| H_i, t_p) = .85$$

30

$$p(\text{not}(\text{Sel_Chart_Pause})|\text{not}(H_i), t_p) = .995$$

35

$$p(\text{Chart_Dialog_Unsucc}, t_o | H_i, t_p) = .08$$

$$p(\text{Chart_Dialog_Unsucc}, t_o | \text{not}(H_i), t_p) = .0005$$

$$p(\text{not}(\text{Chart_Dialog_Unsucc})|\text{not}(H_i), t_p) = .9995$$

40

$$p(\text{not}(\text{Chart_Dialog_Unsucc})| H_i, t_p) = .92$$

45

Inference result : $p(\text{Charting_Assistance}|E, t_p) = .07$

At $t = 70$ seconds

5

$$p(\text{Sel_Chart_Pause}, t_o | H_j, t_p) = .15 \rightarrow .46$$

10

$$p(\text{Sel_Chart_Pause}, t_o | \text{not}(H_j), t_p) = .005 \rightarrow .66$$

$$p(\text{not}(\text{Sel_Chart_Pause}) | \text{not}(H_j), t_p) = .995$$

$$p(\text{not}(\text{Sel_Chart_Pause}) | H_j, t_p) = .85$$

15

$$p(\text{Chart_Dialog_Unsucc}, t_o | H_j, t_p) = .08$$

20

$$p(\text{Chart_Dialog_Unsucc}, t_o | \text{not}(H_j), t_p) = .0005$$

$$p(\text{not}(\text{Chart_Dialog_Unsucc}) | \text{not}(H_j), t_p) = .9995$$

$$p(\text{not}(\text{Chart_Dialog_Unsucc}) | H_j, t_p) = .92$$

25

$$\text{Inference result: } p(\text{Charting_Assistance} | E, t_p) = .04$$

30

At $t = 80$ seconds

35

$$p(\text{Sel_Chart_Pause}, t_o | H_j, t_p) = .15 \rightarrow .46 \rightarrow .85$$

$$p(\text{Sel_Chart_Pause}, t_o | \text{not}(H_j), t_p) = .005 \rightarrow .66 \rightarrow .995$$

40

$$p(\text{not}(\text{Sel_Chart_Pause}) | \text{not}(H_j), t_p) = .995$$

$$p(\text{not}(\text{Sel_Chart_Pause}) | H_j, t_p) = .85$$

45

$$p(\text{Chart_Dialog_Unsucc}, t_o | H_j, t_p) = .08$$

$$p(\text{Chart_Dialog_Unsucc}, t_o | \text{not}(H_i), t_p) = .0005$$

$$5 \quad p(\text{not}(\text{Chart_Dialog_Unsucc}) | \text{not}(H_i), t_p) = .9995$$

$$p(\text{not}(\text{Chart_Dialog_Unsucc}) | H_i, t_p) = .92$$

$$10 \quad \text{Inference result: } p(\text{Charting_Assistance} | E, t_p) = .13$$

At t = 90 seconds

15

$$p(\text{Sel_Chart_Pause}, t_o | H_i, t_p) = .85$$

$$20 \quad p(\text{Sel_Chart_Pause}, t_o | \text{not}(H_i), t_p) = .995$$

$$p(\text{not}(\text{Sel_Chart_Pause}) | \text{not}(H_i), t_p) = .995$$

$$p(\text{not}(\text{Sel_Chart_Pause}) | H_i, t_p) = .85$$

25

$$p(\text{Chart_Dialog_Unsucc}, t_o | H_i, t_p) = .08$$

$$30 \quad p(\text{Chart_Dialog_Unsucc}, t_o | \text{not}(H_i), t_p) = .0005$$

$$p(\text{not}(\text{Chart_Dialog_Unsucc}) | \text{not}(H_i), t_p) = .9995$$

$$35 \quad p(\text{not}(\text{Chart_Dialog_Unsucc}) | H_i, t_p) = .92$$

$$\text{Inference result: } p(\text{Charting_Assistance} | E, t_p) = .04$$

40

System Level IUAF for Multiple Applications

The methods described can not only be applied to individual applications, they can also be employed at the operating system level to reason about providing assistance in a system with multiple applications. In an exemplary embodiment to provide system assistance, a Bayesian network is used to assist a user with exiting an application by guessing when the user has completed or will soon be completing a task.

A Bayesian network knowledge base for assisting a user with this task is displayed in Figure 27. Profile information including whether or not the user has reviewed help in past (Reviewed_Help_in_Past 240), has exited before (User_Exited_Before 241), and has demonstrated that they can perform a double-click operation (Successful_Double_Click 242) are included. Modeled events represented as observations in the Bayesian network include dwell after activity (Dwell after Activity243), sequence of clicks on a sequence of unrelated graphical icons (Click on Seq of Unrelated 244), mouse meandering (Mouse Meandering 245), whether a document has been recently printed (Print Doc 246) or a message has been sent (Send Msg 247), and the duration of time an application has been up (Duration of Time with App 248). A variable represents a probability distribution over the most desired action, allowing us to use probabilistic inference to reason about the likelihood of alternative best actions.

The sample scenario, demonstrates the results of the effect of monitored events on inference about the user. The sample case assumes the following profile: the user has not exited before on his/her own, the user has not demonstrated a successful double-click action, the user reviewed help about exiting applications but did not dwell on the information. The system now detects that the user has been in email for almost 10 minutes, and that the user has clicked on an unrelated sequence of graphical icons. Figure 28 displays the states of each random variable, and shows, by the length of bars next to the states, the probability over different states of each random variable. As demonstrated by the bar graph, the Inference System 76 has computed that there is a high probability the user is ignorant about exiting (User Ignorant about Exit 249) and a

low probability that the user has or will shortly complete a major task (Major Task Completed 250).

In Figure 29, we see the results of inference after new modeled events are considered. In this case, the user has sent email (Send Msg 247) and has now dwelled for more than three scaled seconds after activity (Dwell after Activity 243). Providing a tip to the user is now identified as the most desired exit-related action (Most Desired Exit Action 251). Figure 30 shows the same scenario with one change: the user has received help or assistance within the last two minutes (Recent rec Exit Help 252). Now the most likely desired action is to do nothing, followed by asking the user if the system should exit the application for them (Most Desired Exit Action 251).

IUAF Controller

There is typically a need for the overall control of various components of the Intelligent User Assistance Facility. Controls are needed for making decisions about the policy for calling the Event Processor 74 and Inference System 76, and for determining when assistance should be provided autonomously. Figure 31 is a diagram illustrating various overall coordinative control modes of the Intelligent User Interface controller that may be used separately or together in combination. The first coordinative control mode illustrated is the timer mode 260 in which a repeated system timer fires to notify the IUAF controller to begin another cycle of monitoring and inference analysis (step 261). The IUAF controller then invokes the inference system (step 262). Upon completion of the inference analysis, the controller waits until the timer fires again (step 263). The timer mode can be extended to consider other constraints before allowing a cycle to begin. For example, in cases where complex inference occurs during event monitoring and inference, the additional condition of assessing the availability of computational resources (e.g., identifying whether there is idle time) can be added. Thus, the cycle of event processing and inference will occur when a timer fires and when there is also a condition of computational resource availability.

The second mode illustrated is the demand mode 265. This mode only performs inference analysis when the user asks for assistance. When an instrumented application receives a query for help information (step 266), the instrumented application sends notification to the IUAF controller (step 267). As a result of receiving the notification, the IUAF controller then invokes the inference system (step 268).
5 Upon completion of the inference analysis, the controller waits until it receives the next notification (step 269).

The third mode illustrated is the special event trigger mode 270. In this mode, there is a list of special events that trigger a cycle of inference analysis.
10 Whenever the event processor detects an occurrence of one of the special events (step 271), the event processor sends notification to the IUAF controller (step 272). Again, the IUAF controller then invokes the inference system as a result of receiving the notification (step 273). Upon completion of the inference analysis, the controller waits until it receives the next notification (step 274).

15 Any two or all three modes may be active at the same time. In the case when all three modes are active, if the user has not asked for assistance and has not initiated any user actions that translate into one of the special events during the period of the timer, then a cycle of inference analysis would occur when the period of the timer expired. Otherwise, the cycle of inference analysis would occur during the period of the
20 timer when the user asked for assistance or initiated a user action that translated into one of the special events.

Integration of Analysis of Free-Text Query with Events

As mentioned earlier we can treat words in a user's free-text query for
25 assistance as sources of modeled events. The Microsoft Answer Wizard integrated in all applications of the Microsoft Office '95 product line makes use of a probabilistic model relating words in free-text query to the probability of relevant help topics. This has been described in U.S. Patent Application No. 08400797 and is incorporated hereby by reference in its entirety. In one exemplary embodiment, words and actions can be

handled together in a uniform way as modeled events in a system that reasons about the probabilistic relevance of words and actions.

In an alternate embodiment, there are separate subsystems for handling words and for handling actions. The analyses of these separate subsystems are
5 combined after separate probabilistic analyses are completed. Figure 32 is a flowchart of the method for combining separate probabilistic analyses. First a Bayesian analysis of the most relevant help topics is computed for the given free-text query (step 280). This may be done following the methods described in U.S. Patent Application No. 08400797. The next step is to output a list of help topics in rank order by the highest
10 probability of relevance (step 281). Then a second Bayesian analysis of the most relevant help topics is computed for the given program state, user profile and sequence of user actions (step 282). The next step is to output a second list of help topics in rank order by the highest probability of relevance (step 283). Then both of these lists of help topics are combined, after assigning a weight to the help topics on each list, by
15 computing the probability that each help topic which is the same on both lists is relevant (step 284). The final step is to output the combined list of help topics in order of highest probability of relevance (step 285). In this embodiment, the performance of the system can be tuned by weighting the results of one set of events more heavily than the other by changing the value of a real-numbered parameter that is used to multiply the
20 likelihoods assigned to the probabilities of hypotheses of interest (e.g., relevance of help topics) provided by the word or action analyses.

Building separate systems for analyzing the handling of words in a free-text query and user actions allows the two subsystems to be separately optimized based on specific properties of each class of events. The same can apply to other sets of
25 actions, such as vision and speech events in systems with event monitoring that is extended with these additional sensory modalities.

Separate analyses also allow for the ease of implementing specific policies on combining the results, separating them, or sequencing them. For example, an embodiment might display an initial guess of assistance based on inference based on
30 monitored actions, but when words are added, rely solely on the analysis based on the

words, given the performance of a word analysis systems versus an action system. An alternative approach is to have control rules about the combination versus separation of the inference in different subsystems based on a comparative analysis of the results of the words and actions analysis.

5

Control of Communicating and Performing Assistance

Beyond coordinative control functions, the IUAF controller also makes decisions about when to bother the user with autonomous assistance or with queries for information about the user's goals (based on value of information computations). In an exemplary embodiment, a special Bayesian network, or component of a network, resident in the Inference system 76 computes the overall probability that a user would like assistance at the current time. The IUAF controller uses this changing probability to control when autonomous help is provided. The IUAF provides autonomous help whenever the computed probability that a user would like assistance at that time exceeds a threshold which may be changed by the user. A users changes the threshold for offering autonomous assistance using a threshold controller. The threshold controller, for example, may be in the form of a sliding threshold control can be displayed to users on an autonomous help window, allowing them to change the threshold depending on the value they find in receiving help versus the distraction it causes. A screen shot of a dynamic assistance window with a threshold slider is displayed in Figure 22.

Figure 33 is a flowchart of the method of providing autonomous assistance in this embodiment. During a cycle of inference analysis, the inference system 76 outputs a list of relevant help topics in order of highest probability of relevance (step 290). This may be done as described previously in the discussion of the inference system. The next step is to determine whether the probability of the user needing help exceeds the threshold set by the user for providing autonomous assistance (step 292). If the probability of the user needing help exceeds the threshold set by the user, then the inference system 76 provides a list of relevant help topics in the order of highest probability (step 294). The list may be truncated so as to only show the most

likely or several most likely topics, but within a window that allows the user to scroll through a longer list if desired. Next, the list of relevant topics are checked for redundant help topics that have appeared in the previous list within a specified time period (step 296). If there are no redundant help topics, the list is displayed to the user (step 298). If the user responds to the autonomous offer of help, then steps 290 through 298 are repeated until the user fails to respond before expiration of the response timer. If the user fails to respond before the expiration of the response timer, then a message of apology is displayed on the title of the window (step 299) and the window is subsequently removed.

Those skilled in the art will recognize that alternate embodiments may be used such as only offering autonomous assistance when directly requested by the user, e.g. by clicking on a graphical icon associated with the intelligent assistance facility. The system can also use computed measures of expected utility or its approximation for offering autonomous assistance. Furthermore, the system could only offer autonomous assistance when the sum of the individual probabilities of the relevance of each topic in the set of topics that will be offered to provide assistance exceeds some constant value.

In an exemplary embodiment, guesses on help topics are not presented when a user requests assistance, unless the probability computed for the top five topics together do not sum to at least .40. In yet another alternate embodiment, autonomous assistance will not be offered unless the user is pausing, so as to not distract the user while the user is engaged in activity unless it is deemed important to assist the user with that activity while the activity is in progress.

It can also be useful to provide social conventions and metaphors. For example, the results of inference can be used to control the behavioral gestures of a graphical social cartoon character. An example of this kind of interface is displayed in Figure 34. IUAF inference results about likely problems the user is having or potentially valuable assistance might be alternately be displayed in a "thought cloud" metaphor over a character's head. Free-text queries can be typed into a graphical text input box that appears to be a means for a character to see or listen to the query.

Speech Recognition Enhancement via Integration with IUAF Components

Those skilled in the art of speech recognition will appreciate the power that could be gained if the speech language models used to characterize phonemes by probability were dynamically influenced by information computed about the goals, intentions, or needs of computer users. As portrayed in the influence diagram in Figure 12, user background and goals influence the actions seen in an interface, the words used in free text query, and the speech utterances generated by users in attempting to communicate with a software system or application. Understanding the likelihood of goals and intentions of users in making utterances can enhance our ability to perform automated speech recognition. Such systems will allow us to adjust the probabilities over utterances, and thus, the likelihood of words, and phonemes being uttered by users of computer software.

Speech recognition involves processing the audio signal generated by an utterance into likely phonemes, and the processing of phonemes into likely words. Probabilistic systems for speech recognition may compute the likelihood of phonemes given the audio signal and the likelihood of alternate words given the signal. The term speech is used to refer to the evidence provided by the steps of processing an utterance into audio signals or into phonemes.

Bayesian speech understanding systems compute the probability of word or stream of word from an utterance as follows:

$$p(\text{words}|\text{speech}) = \frac{p(\text{words}) * p(\text{speech} | \text{words})}{p(\text{speech})}$$

By computing the probability over a user's needs or goals given user actions as represented by modeled events (including those based on gaze), we can refine the accuracy of the probability of words given a user speech as follows:

We can update the likelihood that words or word strings will be used in the context of goals by representing the likelihood that words or word strings will be used in the context of specific goals, and update the base probabilities that words will be used, $p(\text{words})$, contained in a typical language model in real-time. That is, instead of

using a static database of $p(\text{words})$, we instead update the probability of words as follows:

$$p(\text{words}) = \sum_i p(\text{words} | \text{user needs}_i) * p(\text{user needs}_i)$$

- 5 This computation takes advantage of a knowledge base of context-specific language models that contain information of the form $p(\text{words}) = \sum_i p(\text{words} | \text{user needs}_i)$.

At run-time, actions are analyzed by the Event Monitoring System and are passed to the Inference System which provides dynamically the continually recomputed probability distribution over user needs, $p(\text{user needs} | \text{events})$ as described
 10 above. One exemplary application is the reasoning about help topics, and context specific language models contain information of the form $p(\text{words} | \text{relevant help topic})$.

In a more general statement showing the modification of the above equation for speech recognition, we input dynamic events

15

$$p(\text{words} | \text{speech}, \text{events}) =$$

$$\left[\sum_i p(\text{words} | \text{user needs}_i, \text{events}) * p(\text{user needs}_i | \text{events}) \right] * p(\text{speech} | \text{words})$$

20

$$p(\text{speech})$$

This can be reduced to:

$$25 \quad p(\text{words} | \text{speech}, \text{events}) = \frac{p(\text{words} | \text{events}) * p(\text{speech} | \text{words})}{p(\text{speech})}$$

Figure 23 displays the key components and flow of analysis employing
 30 key components of the intelligent user assistance facility to enhance speech recognition by dynamically changing the likelihoods over words, and thus, phonemes uttered. As displayed in the figure, user actions are tracked by an instrumented program 72, and atomic events are composed into modeled events by the Event Processor 74 and passed to the Inference System 76. The inference system computes a probability distribution of

the needs of a user and the weighted list of needs are passed to a Dynamic Language Model Generator 300.

The Dynamic Language Model Generator 300 has a mapping between alternative needs and the likelihood of utterances or words. This component may contain information from a Language Model Database 301 that contains language models containing probabilistic information of the form, $p(\text{words} \mid \text{user need})$, the probabilities of different words or word strings being uttered by a user given the user needs. These language models can be constructed through training with statistics from a large corpus of strings of words that are recorded in different user-needs contexts. At run-time, the Inference System provides a probability distribution over user needs which is used to control the creation of an updated language model, based on weighting the language models from the different contexts. The updated language model 302 contains information of the form, $p(\text{words} \mid \text{events})$, created by weighting sets of words in terms of the likelihood of alternative needs based on the probabilistic inference results.

Overall, taking advantage of context in listening can enhance the accuracy of speech recognition system 303 by lowering the probability of utterances that are unlikely and raising the probability of utterances that are more likely, per the user modeling.


Beyond use with software in receiving assistance, the dynamic speech understanding has application in variety of speech understanding tasks. As one example, the methods have application in a variety of interfaces with command and control, referring to the manipulation of controls and functionality of software. If we model the goals of the user, we can generate grammars that are appropriately weighted to reflect the next steps of the user, and there are less prone to error.

These methods can be used in conjunction with models that generate the probability of the next software command given user goals. In this approach we focus on a particular user need, the desire for action. The user's desire for action may refer to the implementation of one or more software controls by speech, including the access of one or more help topics. We substitute $p(\text{desired action} \mid \text{events})$ for $p(\text{user needs} \mid$

events) in the above equations, and employ language models that contain information $p(\text{words} \mid \text{desired action})$, instead of $p(\text{words} \mid \text{user needs})$. Events considered in this model can include context which may include information about the window that is active and the text, buttons, and controls that are displayed to the user while the speech
5 utterance is occurring.

For example, dynamic language models for command and control have application in systems that explicitly generate and display probabilistic results as part of their functioning. Consider as an example, a probabilistic reasoning system that employs a Bayesian network or influence diagram for assisting a physician with
10 diagnosis from a set of symptoms. Bayesian reasoning systems in medicine may have sets of windows containing the results of diagnostic inference. A window may list diseases by computed likelihoods, another may list the best next symptoms or tests to evaluate, based on value of information. By building or assessing a model of the probability of the user's actions given the context of displayed information, which may
15 include functions of the computed probabilities and expected utilities assigned to displayed objects, we can compute and use a dynamically re-computed $p(\text{words} \mid \text{events})$ in speech recognition.

While the present invention has been illustrated with reference to exemplary embodiments, those skilled in the art will appreciate that various changes in
20 form and detail may be made without departing from the intended scope of the present invention as defined in the appended claims. For example, the general event composition and inference system may also be used in an operating system to provide an intelligent user shell and, more generally, to optimize the functionality of any computer system or software. Because of the variations that can be applied to the
25 illustrated and described embodiments of the invention, the invention should be defined solely with reference to the appended claims.



Appendix 1. More complete Bayesian network for assisting users with a spreadsheet program

Claims

1. A method of building an intelligent user interface facility for a software program to assist a user in operating the software program,
comprising the steps of:
identifying program functionality of the software program for providing intelligent user assistance,
building a reasoning model for the task of providing the intelligent user assistance for the identified program functionality of the software program,
defining modeled events for the identified program functionality for use by the reasoning model to provide intelligent user assistance, and
identifying for each modeled event one or more atomic user-interface actions that signify the occurrence of each modeled event.
2. The method of claim 1 further comprising the steps of:
constructing an event monitor to capture the identified atomic user-interface actions, and
integrating the event monitor with the software program to instrument that program for intelligent user assistance.
3. The method of claim 1 further comprising the step of:
constructing an event processor to compose modeled events from atomic user-interface actions.
4. The method of claim 1 further comprising the step of:
constructing a knowledge base from the model, and
constructing an inference engine from the model which uses the knowledge base to determine suggested user assistance for the occurrence of modeled events.

5. The method of claim 4 wherein the determination of suggested user assistance is performed by a probability analysis inferring the most likely user assistance that is helpful given the occurrence of specific modeled events.

6. The method of claim 5 wherein a Bayesian network is employed to perform the probability analysis.

7. The method of claim 5 wherein a Bayesian influence diagram is employed to perform the probability analysis.

8. The method of claim 1 wherein the modeled events include contextual information about the state of data structures in the program.

9. The method of claim 1 wherein the modeled events include words in a free-text query seeking help information in using the program.

10. The method of claim 1 wherein the software program executes on a computer system with a speech recognition system and the modeled events include words identified by the speech recognition system.

11. The method of claim 1 wherein the software program executes on a computer system, that includes a display screen and a vision recognition system, and the modeled events include the focus, identified by the vision recognition system, of a user's gaze on the screen.

12. The method of claim 1 wherein the modeled events include menu commands of the program.

13. The method of claim 1 wherein the modeled events include a time period during which there is no user activity.

14. The method of claim 1 wherein the software program executes on a computer system with a display screen and the modeled events include text displayed on the screen followed by a period of dwelling and a subsequent scroll.

15. The method of claim 14 wherein the text displayed on the screen is help information and the modeled event is used for indicating a user's competency in using specific program functionality.

16. The method of claim 14 wherein the modeled event of the text displayed on the screen followed by a period of dwelling and a subsequent scrolls is used for tracking a user's indication of interest in the textual information.

17. The method of claim 16 wherein the text is stored in the modeled event database.

18. The method of the claim 16 wherein the path of the text's location is stored in the modeled event database.

19. A computer readable medium having computer executable instructions for performing the steps of claim 1.

20. An intelligent user interface facility for a software program to assist a user in operating the software program,

comprising:

an event processor that obtains events from the software program, and

an inference system that determines suggested user assistance from the occurrence of the events.

21. The intelligent user interface facility of claim 20, wherein the software program is instrumented to monitor user interface actions and to extract user interface events.

22. The intelligent user interface facility of claim 20, wherein the software program is an application program.

23. The intelligent user interface facility of claim 20, wherein the software program is an operating system program.

24. A computer readable medium having computer executable components of claim 20.

25. The intelligent user interface facility of claim 20, wherein the event processor composes the events into modeled events.

26. The intelligent user interface facility of claim 20, further comprising:
a user profile system that maintains an historical record of events for the user that is used by the inference system in determining suggested user assistance from the occurrence of events.

27. The intelligent user interface facility of claim 26, wherein the modeled event database is initialized by persistent modeled events for the program which are stored in the user profile of the user of the program.

28. A computer readable medium having computer executable components of claim 26.

29. The intelligent user interface facility of claim 20, further comprising:
a controller for initiating a cycle of inference analysis by the inference system.

30. The intelligent user interface facility of claim 29, wherein the controller initiates a cycle of inference analysis on a periodic cycle.

31. The intelligent user interface facility of claim 29, wherein the controller initiates a cycle of inference analysis whenever one or more events has occurred.

32. The intelligent user interface facility of claim 29, wherein the controller initiates a cycle of inference analysis whenever the user activates the help facility of the program.

33. The intelligent user interface facility of claim 20, wherein the suggested user assistance is a list of help topics.

34. The intelligent user interface facility of claim 33, wherein the inference system uses probabilistic inference analysis to determine the list of help topics and the inference system lists only those help topics which have a likelihood of relevance given the occurrence of the events that exceed a confidence measure.

35. The intelligent user interface facility of claim 34, wherein the inference system lists those help topics, which have a likelihood of relevance that exceed the confidence measure, in a scrollable dialog box in order of highest probability value.

36. The intelligent user interface facility of claim 34, wherein the inference system

offers assistance autonomously to the user whenever the likelihood that a help topic is relevant given the occurrence of the events exceeds a threshold value.

37. The intelligent user interface facility of claim 36, wherein the threshold value is set by the user.

38. The intelligent user interface facility of claim 34, wherein the inference system combines the probabilistic inference analysis to determine the list of help topics given the occurrence of events with a probabilistic inference analysis to determine a list of help topics given a free-text query.

39. The intelligent user interface facility of claim 38, wherein the inference system lists only those help topics of the combined analysis which have a likelihood of relevance that exceed a confidence measure.

40. The intelligent user interface facility of claim 34, wherein the inference system only uses a probabilistic inference analysis of the words in the free-text query to determine a list of help topics whenever a user enters a free-text query.

41. The intelligent user interface facility of claim 20, wherein the inference system determines suggested user assistance using probability and offers assistance autonomously to the user whenever the inference system determines that the probability that the user needs assistance given the occurrence of the events exceeds a threshold value.

42. The intelligent user interface facility of claim 41, wherein the threshold value is set by the user.

43. The intelligent user interface facility of claim 20, wherein the inference system requests information autonomously from the user whenever the inference system computes that the value of acquiring information from the user is greater than a threshold value.

44. The intelligent user interface facility of claim 43, wherein the threshold value is set by the user.

45. The intelligent user interface facility of claim 20 wherein the inference system determines suggested user assistance from both the occurrence of the events and a free-text query submitted by the user.

46. The intelligent user interface facility of claim 20 wherein the suggested user assistance is provided to the user through a social user interface in the form of an animated character.

47. The intelligent user interface facility of claim 20, wherein the event processor and the inference system are implemented as a single component chosen from the group consisting of a software program, a software library or an instantiated object.

48. The intelligent user interface facility of claim 20, wherein the event processor obtains only a specific number of the last occurring events during a cycle of inference analysis.

49. An event processor in a computing device having an event source comprising:

- a local store for temporarily storing the events;
- a composer for generating modeled events from one or more events; and
- a database for storing the modeled events generated.

50. The event processor of claim 49, wherein the events stored in the local store are assigned a horizon.

51. The event processor of claim 49, wherein the modeled events generated have an assigned horizon.

52. The event processor of claim 49, wherein the events are user-interface actions.

53. The event processor of claim 49, wherein the source of events is chosen from the group consisting of a software program, a software library or an instantiated object.

54. The event processor of claim 49, wherein the events are atomic events and the composer generates modeled events from one or more atomic events.

55. The event processor of claim 49 wherein the events are atomic events and the composer generates modeled events from one or more atomic events and from one or more modeled events.

56. The event process or claim 49 wherein the composer generates modeled events from one or more modeled events.

57. The event processor of claim 49 wherein the local store is a circular queue.

58. The event processor of claim 49 wherein the local store is a database.

59. The event processor of claim 49 wherein the database for storing the modeled events is in system cache.

60. The event processor of claim 49, wherein there is a single record for every modeled event for that database.

61. The event processor of claim 49 wherein the database contains records that include a set of counters of the number of times a modeled event has occurred.

62. The event processor of claim 49 wherein a record of the database contains horizon and probability dynamics information.

63. The event processor of claim 49 wherein a record of the database is marked for persistent storage.

64. The event processor of claim 49, wherein the modeled events generated are logical events.

65. The event processor of claim 49, wherein there is a separate database for each source of events that stores the modeled events generated.

66. The event processor of claim 49, wherein there is a separate database for each source of events that stores the modeled events generated and a general database that stores common system or application events.

67. A computer readable medium having computer executable components of claim 49.

68. An event system specification tool for a computing device comprising:
an event language for specifying modeled events; and
an event language interpreter for translating statements of the event language into a target code language.

69. The event system specification tool of claim 68, wherein the event language interpreter is used for translating statements of the event language into modeled database definitions.

70. The event system specification tool of claim 68, wherein the event language includes temporal operators.

71. The event system specification tool of claim 68, wherein the event language includes Boolean operators.

72. The event system specification tool of claim 68, wherein the event language includes set-theoretic operators.

73. The event system specification tool of claim 68, wherein the modeled events are some of the variables of a reasoning model and the event language includes temporal functions that define the persistence of probabilistic relationships between one or more modeled events.

74. The event system specification tool of claim 68, wherein the modeled events are some of the variables of a reasoning model and the event language includes temporal functions that define the persistence of probabilistic relationships between one or more modeled events and one or more other variables in the reasoning model.

75. The event system specification tool of claim 74, wherein the reasoning model is a Bayesian network.

76. The event system specification tool of claim 74, wherein the reasoning model is rule-based.

77. A computer readable medium having computer executable components of claim 68.

78. The event system specification tool of claim 70, wherein a temporal operator is chosen from the group consisting of rate, sequence and dwell.

79. An inference system for an intelligent user interface comprising:

a knowledge base that includes temporal reasoning knowledge; and
an inference engine that includes temporal reasoning component.

80. The inference system of claim 79 wherein the knowledge base is a Bayesian network.

81. The inference system of claim 79 wherein the knowledge base is a set of logical rules.

82. The inference system of claim 79 wherein the inference engine is a Bayesian network.

83. The inference system of claim 79 wherein the inference engine further includes a value-of-information component that computes the probability that the value of acquiring information from the user is greater than a threshold value.

84. A computer readable medium having computer executable components of claim 79.

85. A Bayesian network comprising:
nodes representing hypotheses,
nodes representing observed evidence, and
temporal reasoning approximation function associated with at least one of the observed evidence nodes to compute the probability of the relevance as it changes with the passage of time.

86. The Bayesian network of claim 85, wherein the observed evidence are modeled events of user actions.

87. A method for an inference system to autonomously acquiring unobserved information, given the occurrence of specific events, from a user to determine suggested user assistance, comprising computing the probability that the value of acquiring information from the user is greater than a threshold value.

88. A user profile system for providing intelligent user-interface assistance for a software program, comprising:

- a user profile database which persistently stores records of user's experience with functionality of the software program,

- a user profile access component that access a user's records of experience with functionality of the software program , and

- an event status update component which updates the user's records of experience with functionality of the software program as the user interacts with the software program.

89. The user profile system of claim 88, further comprising:

- a user competency dialogue component which establishes user competency by obtaining information about a user's background with the functionality of the program.

90. The user profile system of claim 88, further comprising:

- a user competency dialogue component which establishes user competency by obtaining information about a user's goals with the functionality of the program.

91. The user profile system of claim 88, further comprising:

- a user competency dialogue component which establishes user competency by obtaining information about a user's experience with the functionality of the program.

92. The user profile system of claim 88, further comprising:

- a customize tutorial component which list in rank order the most useful information topics to train the user in the functionality of the software program.

93. The user profile system of claim 88, wherein the user profile database contains at least one record of a successful completion of a program command.

94. The user profile system of claim 88, wherein the user profile database contains at least one record of a unsuccessful completion of a program command.

95. A computer readable medium having the computer executable components of claim 88.

96. An intelligent user interface facility for a software program to assist a user in operating the software program running on a computing device with a speech recognition system having a language model with probabilities associated with words of the language,

comprising:

an event processor that obtains events from the software program,

an inference system that determines suggested user assistance from the occurrence of the events,

a dynamic language model generator that uses the suggested user assistance determined by the inference system to modify probabilities of the words of the language recognized by the speech recognition system .

97. The intelligent user interface facility of claim 96, wherein the software program is instrumented to monitor user interface actions and to extract user interface events.

98. The intelligent user interface facility of claim 97, wherein the software program is an application program.

99. The intelligent user interface facility of claim 97, wherein the software program is an operating system program.

100. The intelligent user interface facility of claim 99, wherein the operating system uses the speech recognition system for command and control of the computing device.

101. A computer readable medium having the computer executable components of claim 96.

102. A method for modifying the probabilities of words in a language model database that are recognized by a speech recognition system comprising the steps of:

receiving user actions from an instrumented program,
composing modeled events from the user actions,
employing probabilistic inference to compute a probability distribution of the needs of a user from the modeled events, and
modifying the probabilities of the words in the language model database that are recognized by the speech recognition system using the probability distribution.

103. The method of claim 102, wherein the instrumented program is used for directing software controls with speech.

104. The method of claim 102, wherein the instrumented program is used for receiving spoken help queries.

105. A computer readable medium having computer executable instructions for performing the steps of claim 102.

1 / 3 5

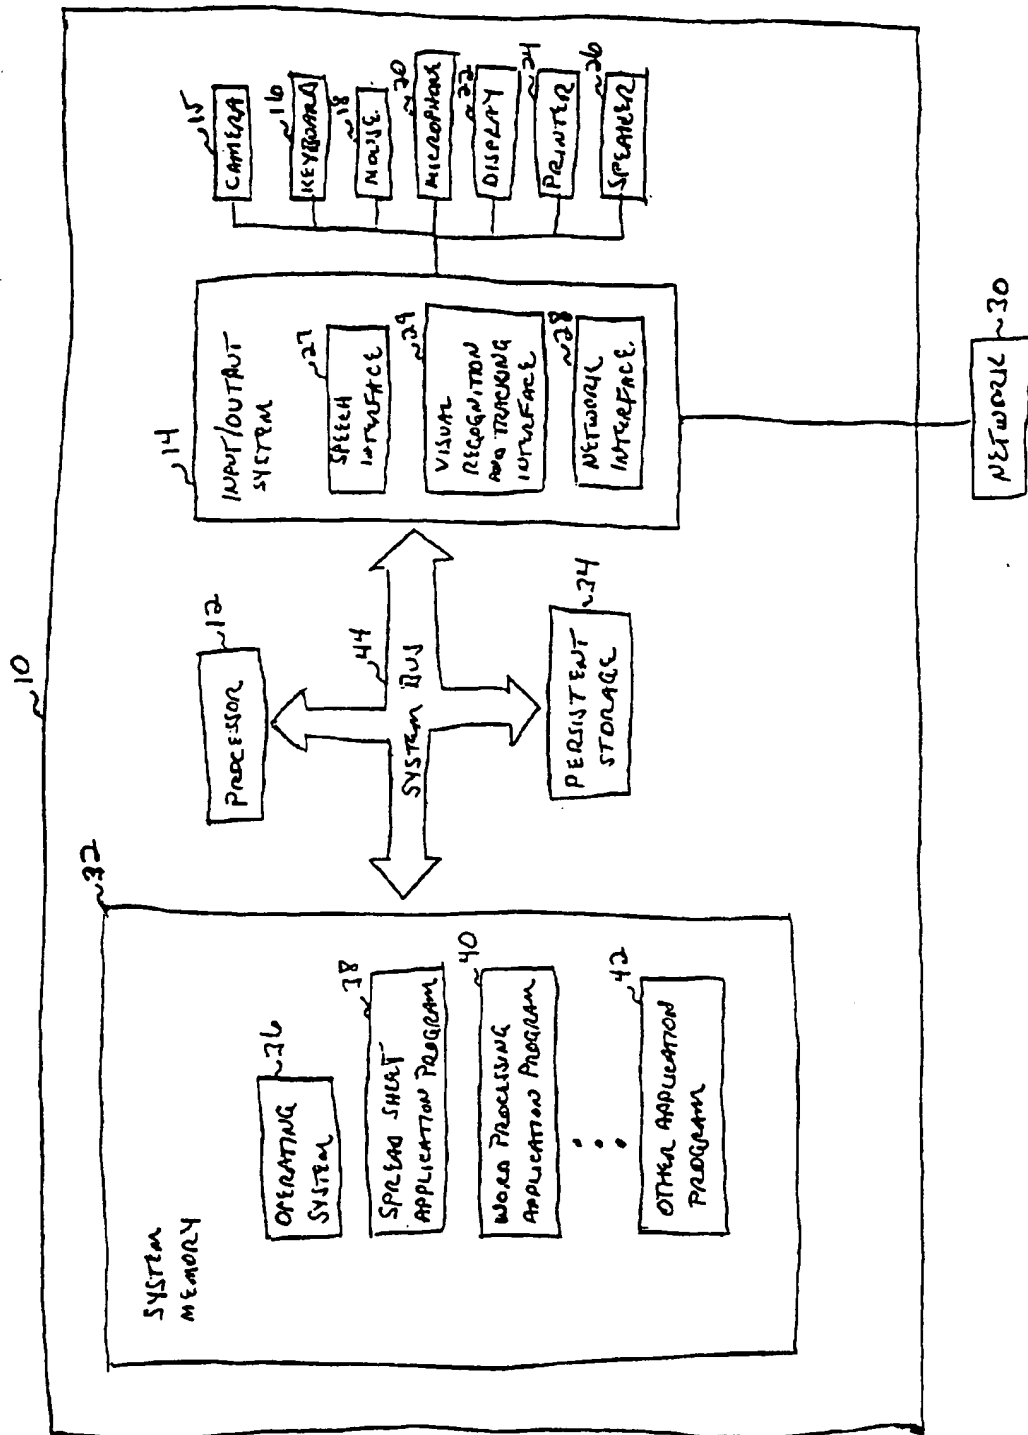


FIGURE 1 BLOCK DIAGRAM OF COMPUTER SYSTEM

2 / 3 5

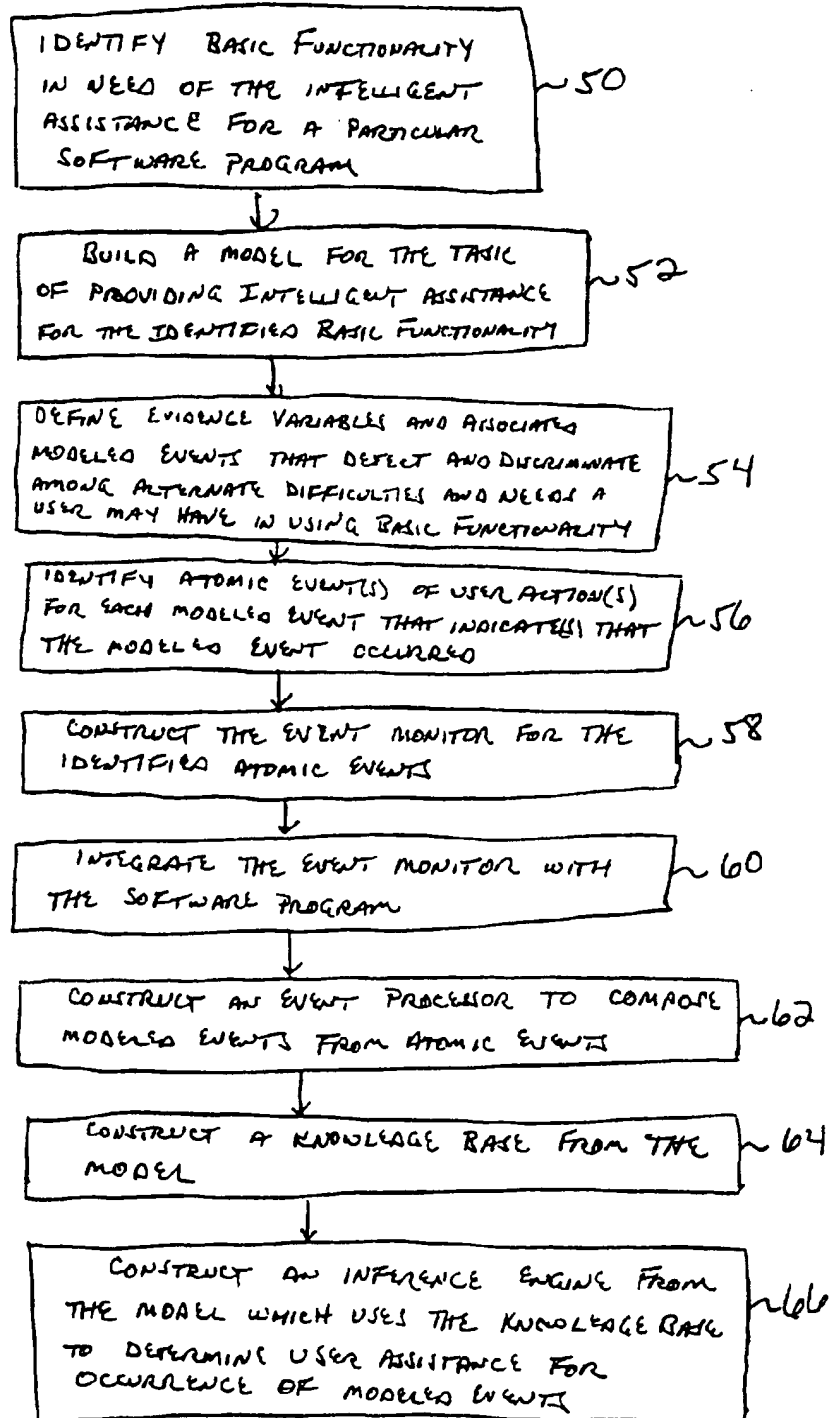


FIGURE 2

METHOD FOR MAKING A GUI FOR A PROGRAM

3 / 3 5




Figure 3. Diagram illustrating a simple, two-variable Bayesian network

4 / 3 5




FIG 4 BLOCK DIAGRAM OF INTELLIGENT USER ASSISTANCE FACILITY

5 / 3 5




FIGURE 5
BLOCK DIAGRAM OF INSTRUMENTED
APPLICATION PROGRAM

6 / 3 5

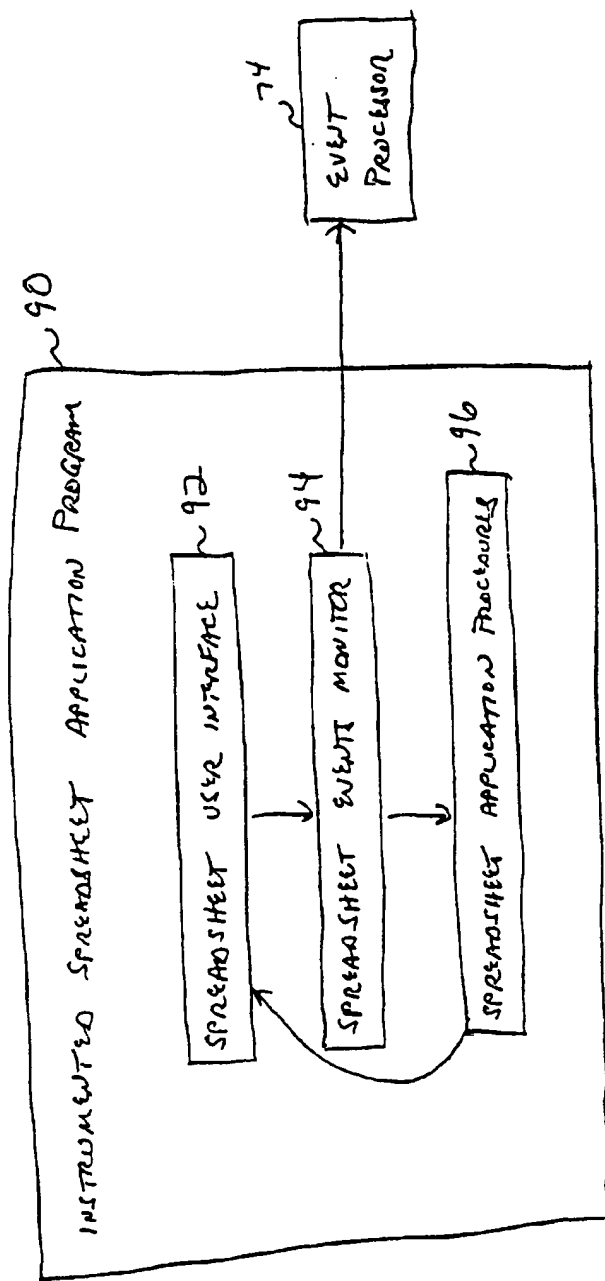


FIGURE 6
BLOCK DIAGRAM OF AN
INSTRUMENTED SPREADSHEET APPLICATION

7 / 3 5




FIGURE 7 BLOCK DIAGRAM OF EVENT PROCESSOR

8 / 35

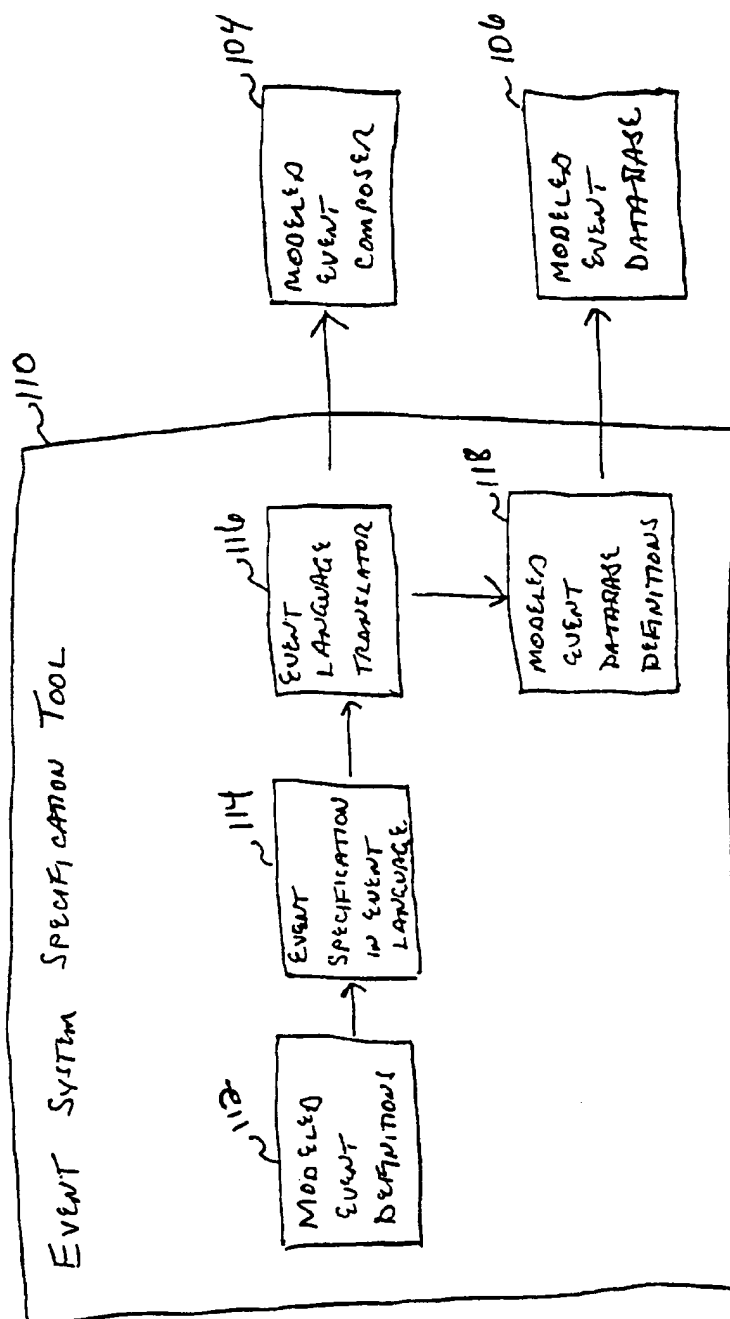
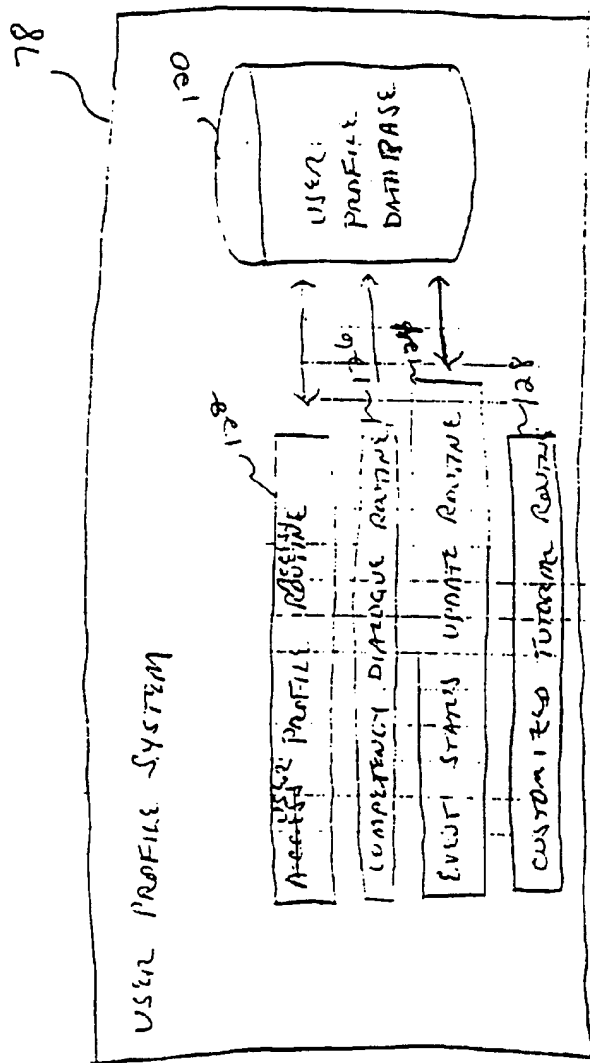


FIGURE 8

9 / 3 5



Block Diagram
of
USER PROFILE SYSTEM

FIG. 9

10/35

when do
records get
deleted from user
profile? time based?




FIG 10

Flowchart of Event
Status Update Routine

11 / 35

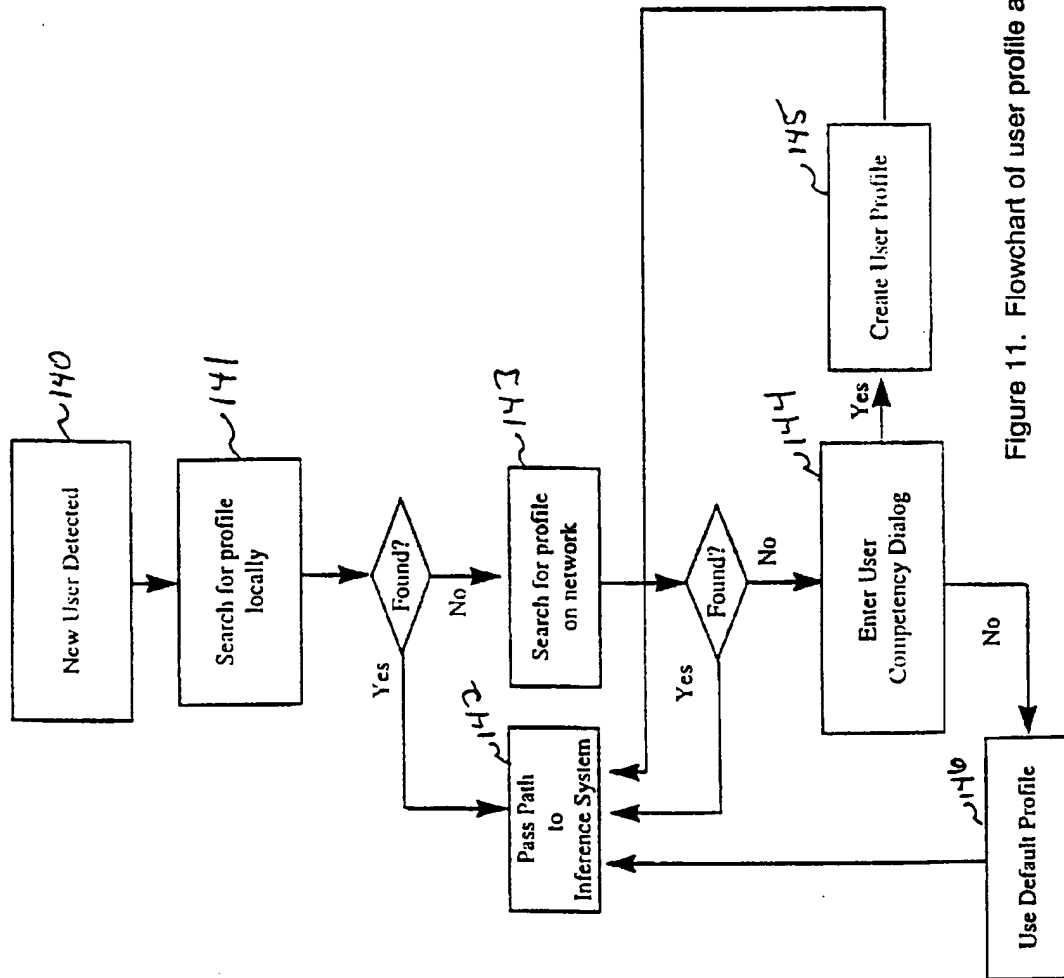


Figure 11. Flowchart of user profile access routine.

12/35

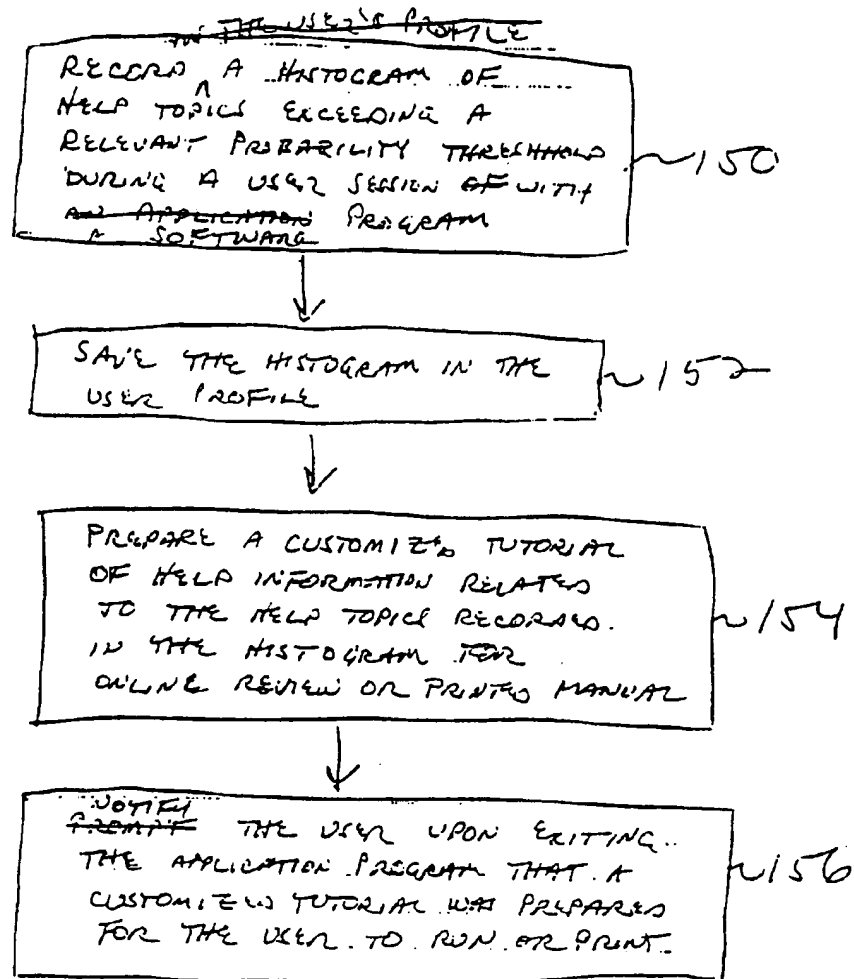
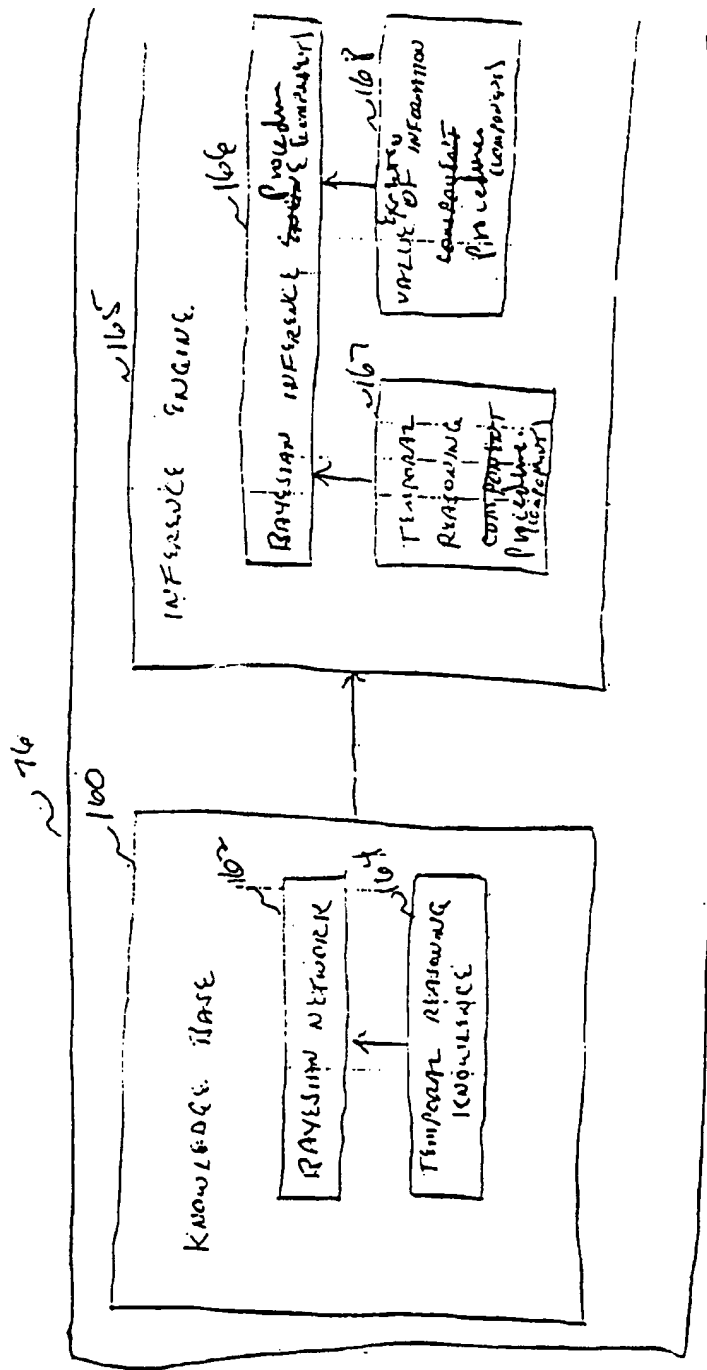


FIG 12

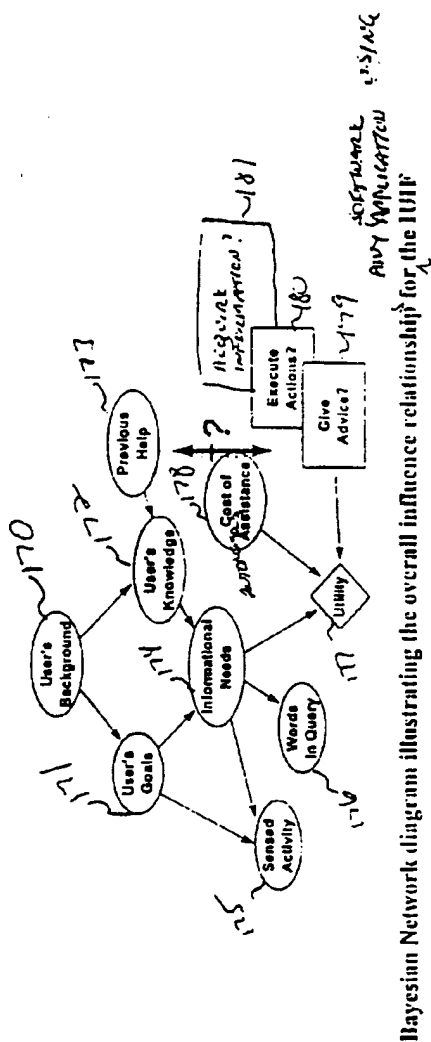
FLOWCHART FOR GENERATION
OF CUSTOMIZED TUTORIAL

13 / 35



Block Diagram of

FIG. 13



Bayesian Network diagram illustrating the overall influence relationship⁵ for the IUIP.

F1a

181




Figure 15. A portion a Bayesian network representing some probabilistic dependencies over time.

16 / 35

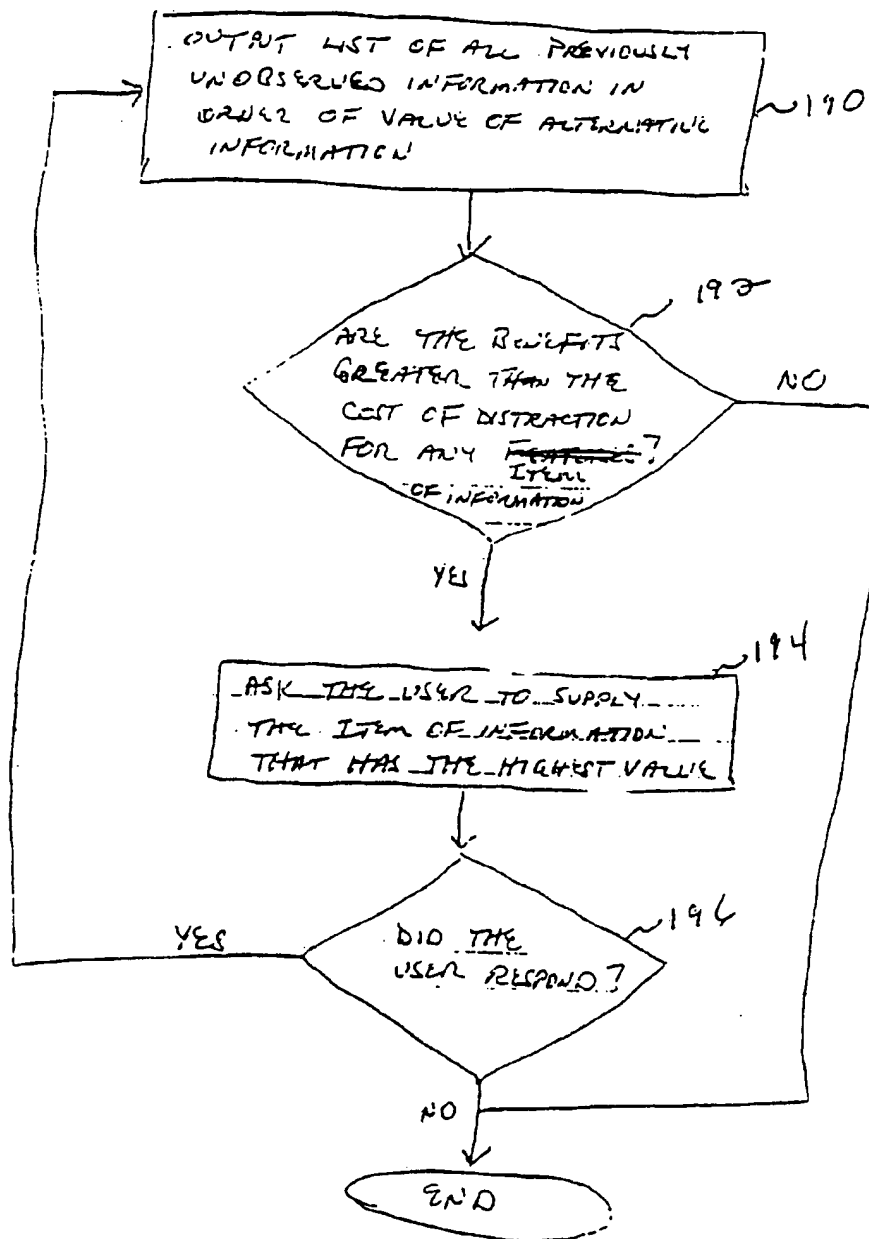


FIG. 11c

FLOWCHART OF
VALUE-OF-INFORMATION
PROCEDURE

17 / 3 5

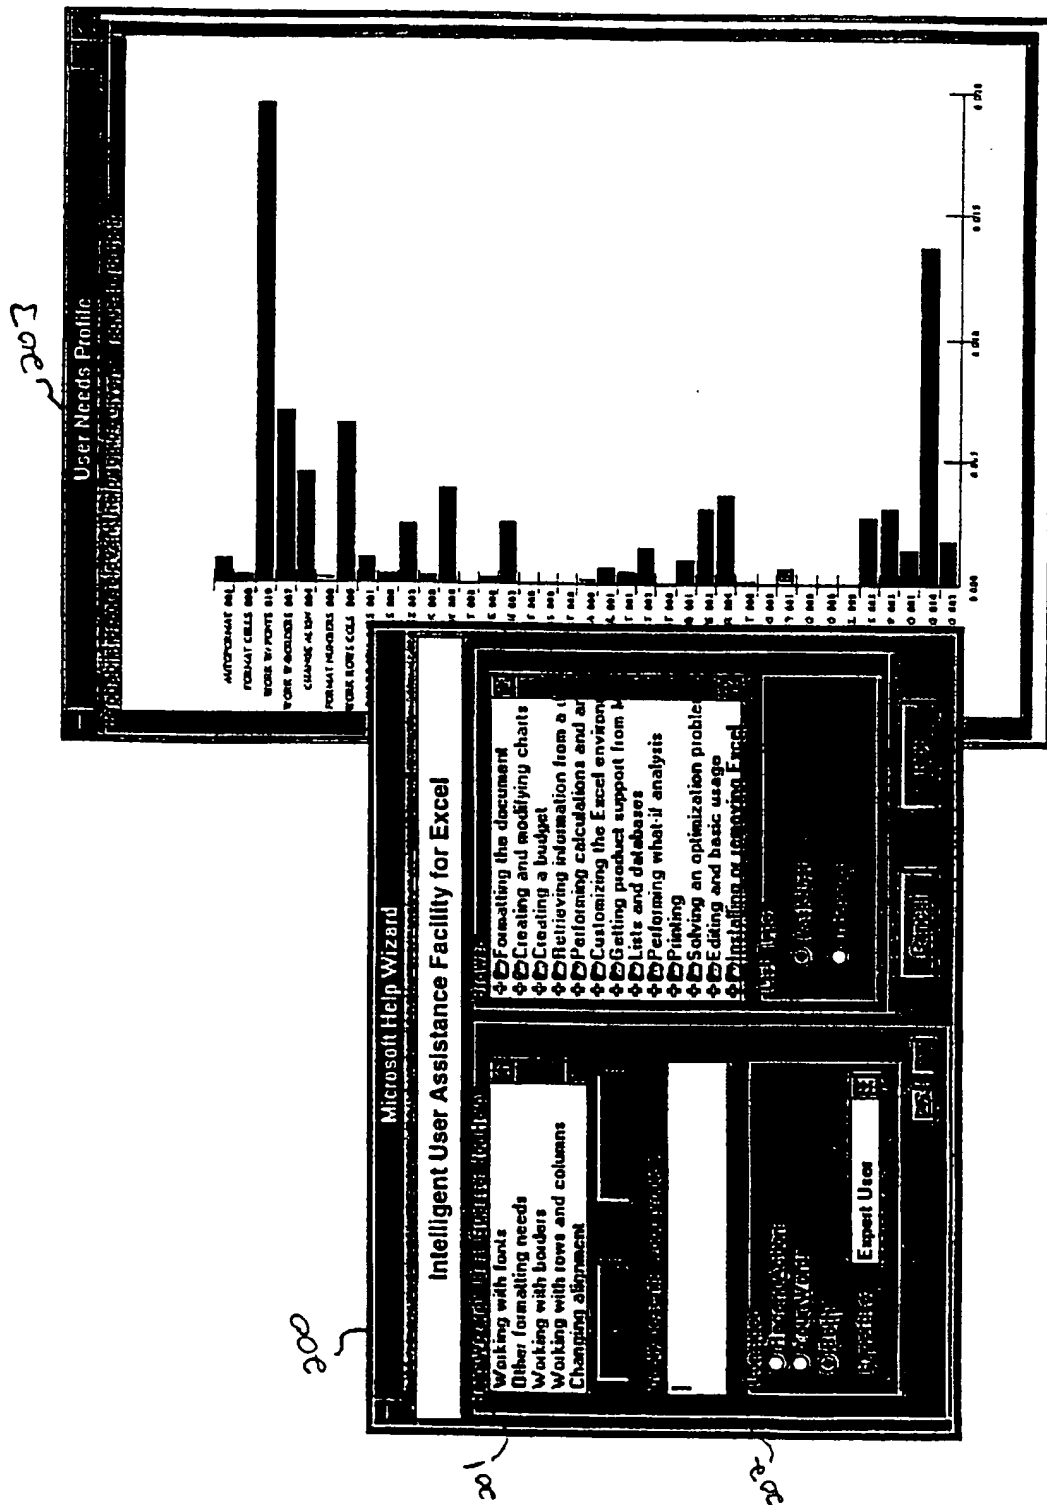


Figure 17. Screen outputs of the IUAF in an exemplary embodiment of a spreadsheet program.

18 / 3 5

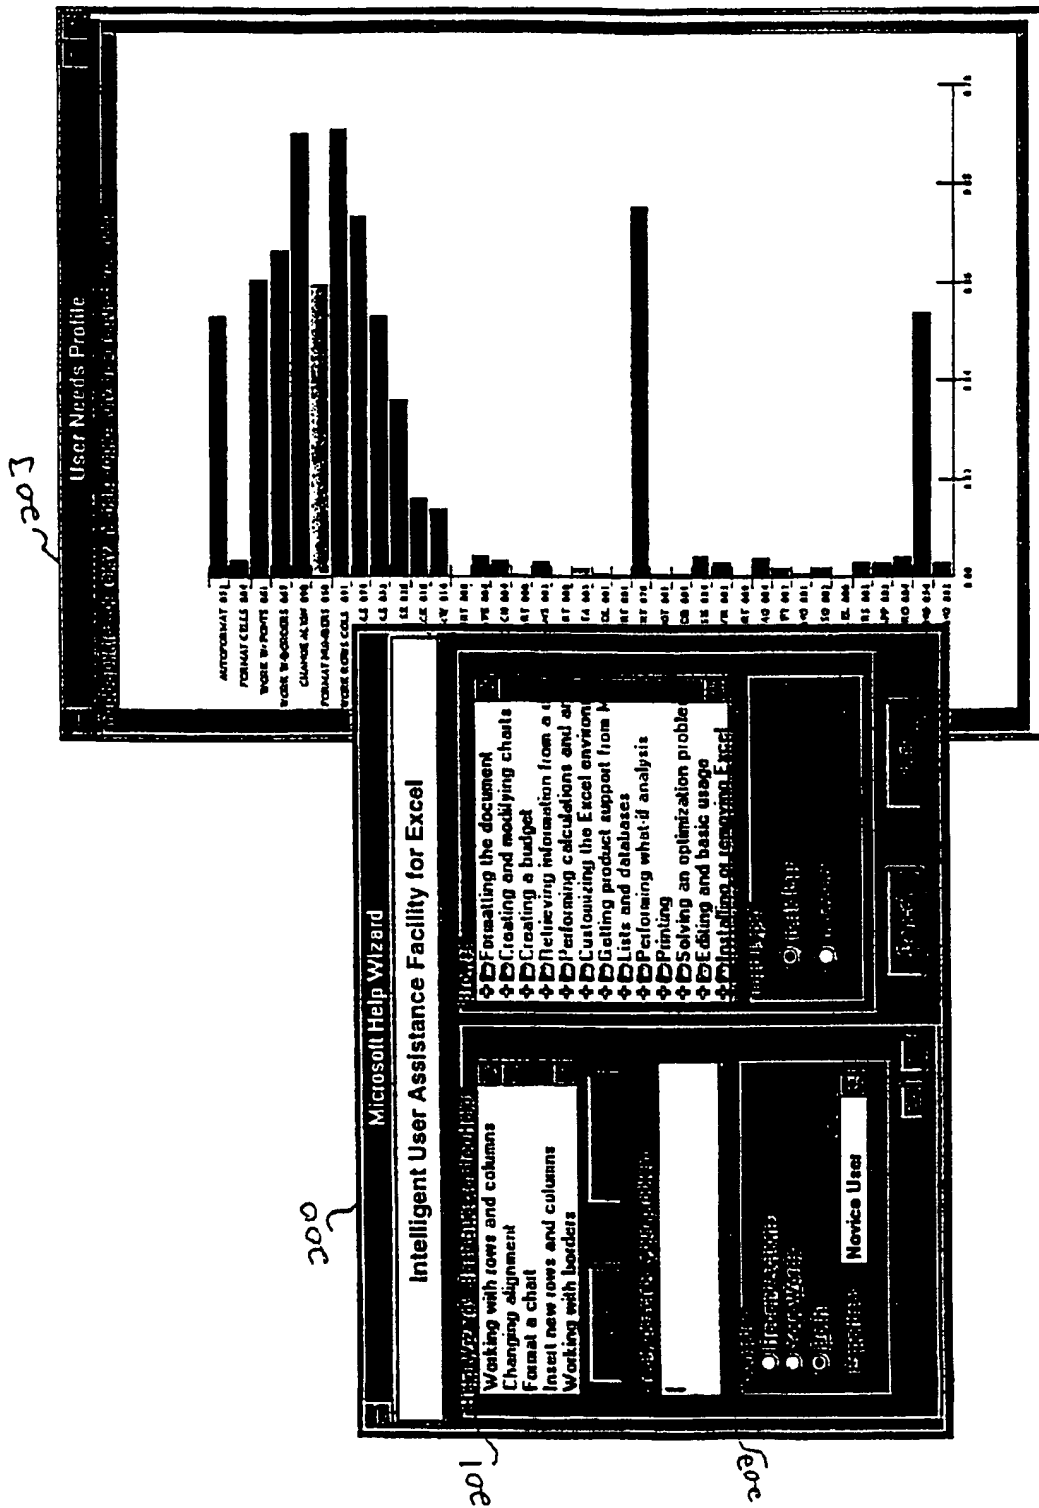


Figure 18. Screen outputs of the IUAF in an exemplary embodiment of a spreadsheet program.

19 / 3 5

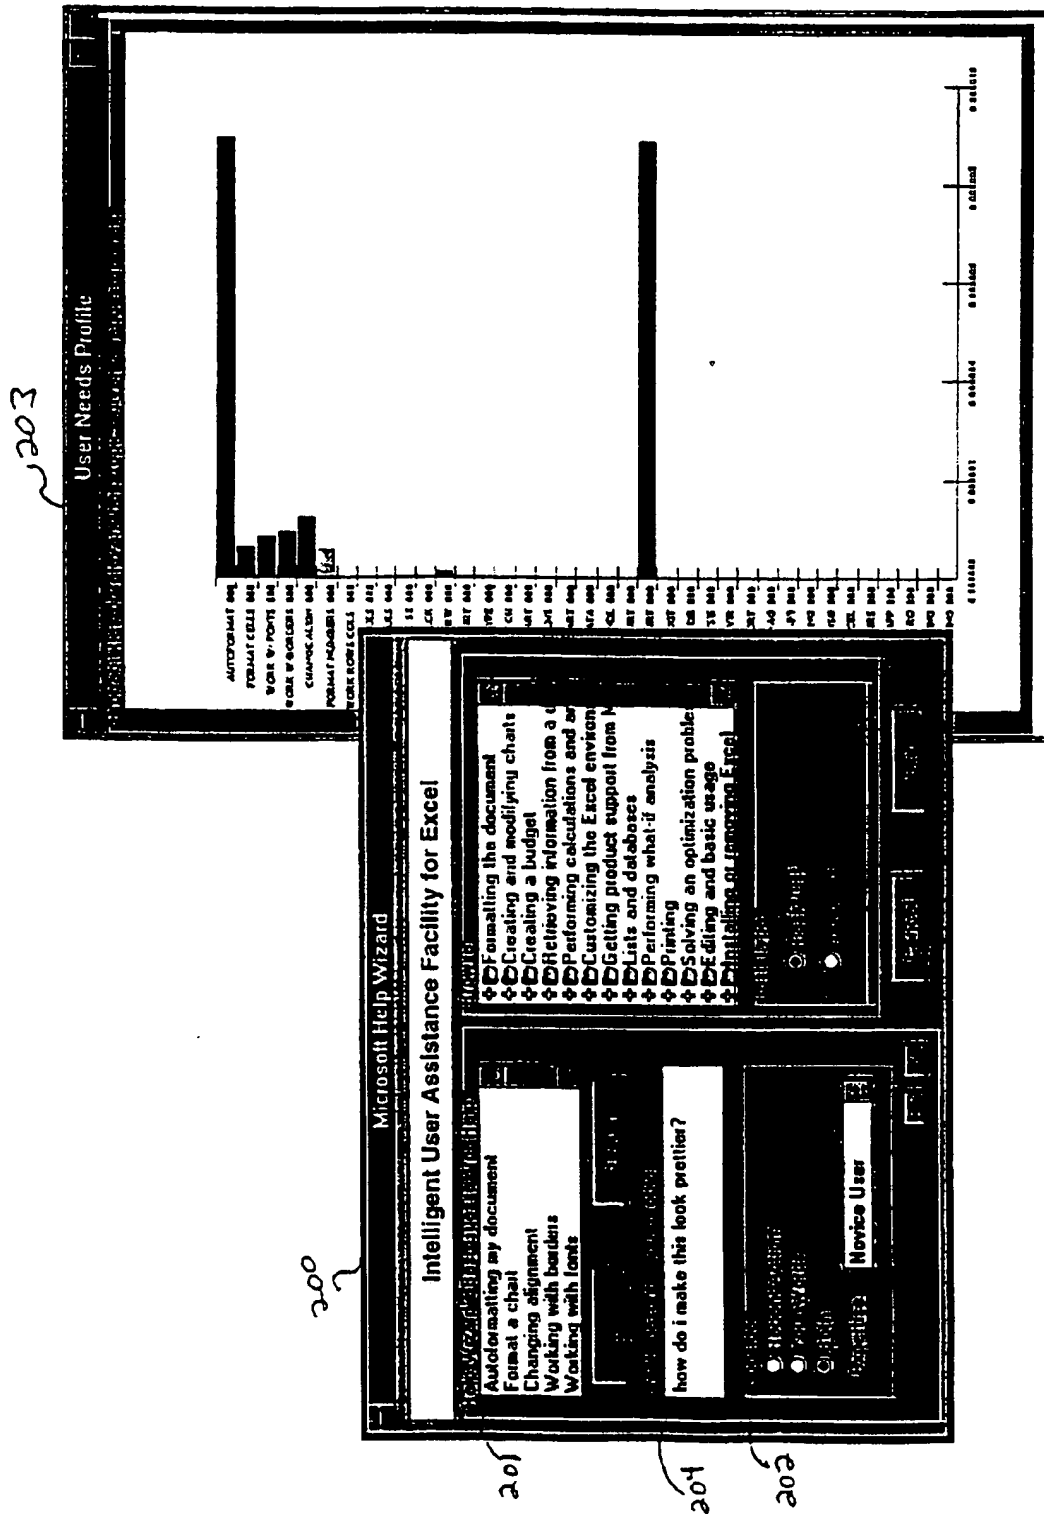


Figure 19. Screen outputs of the IUAF in an exemplary embodiment of a spreadsheet program.

20 / 35

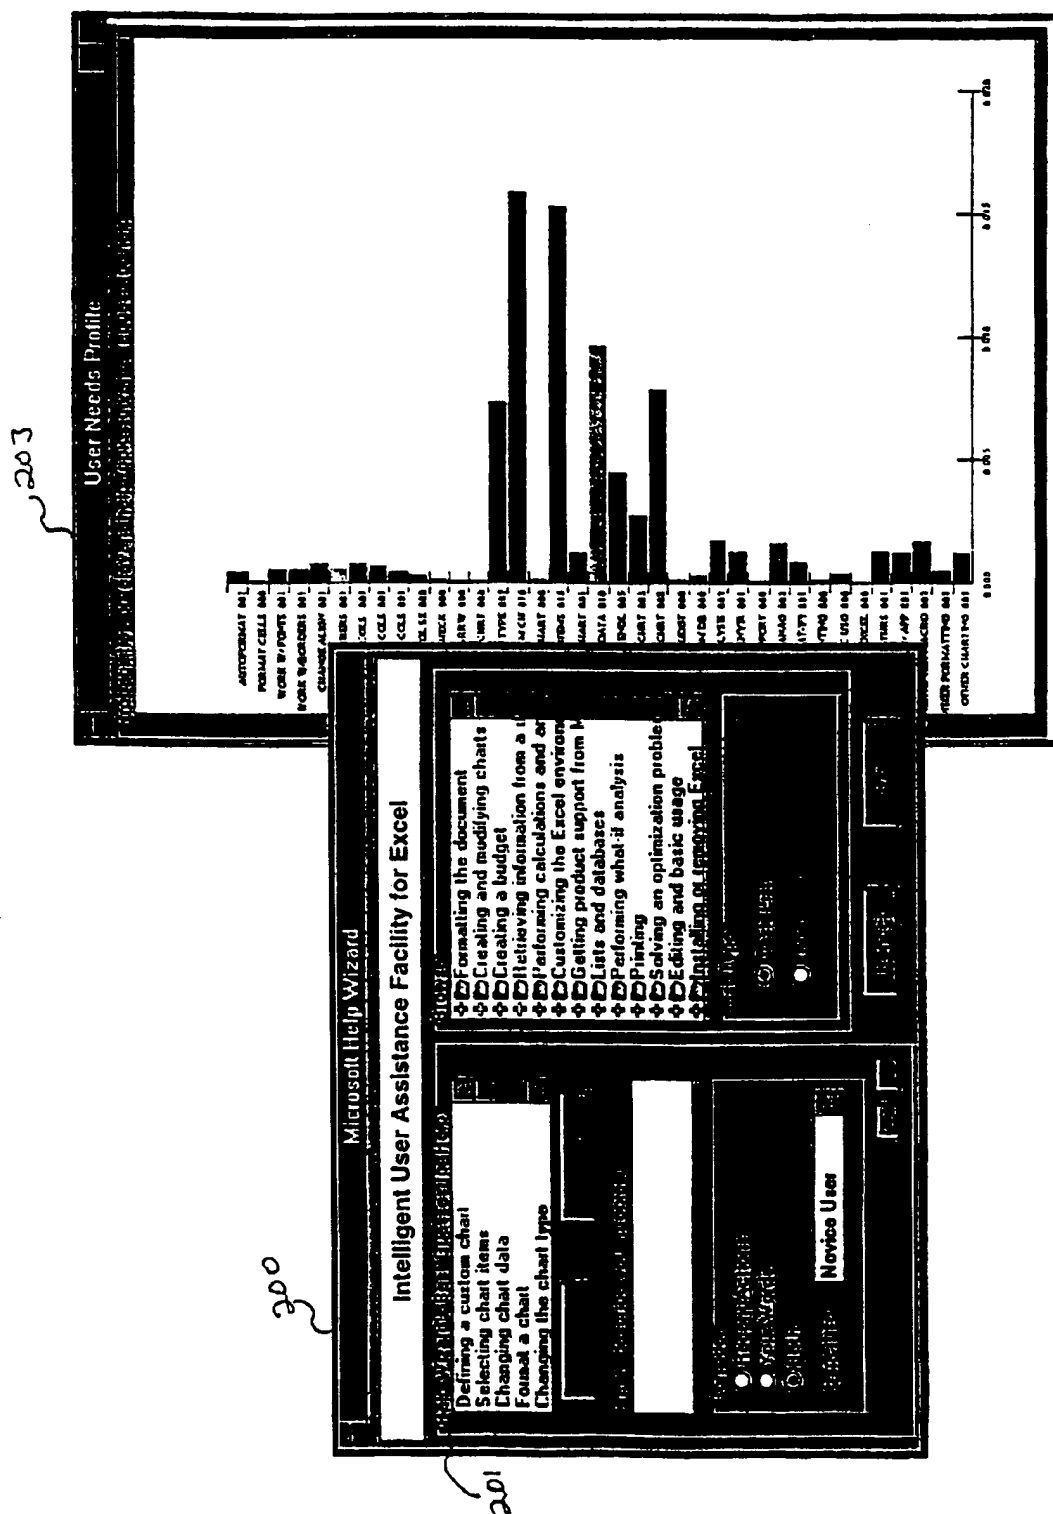


Figure 20. Screen outputs of the IUAF in an exemplary embodiment of a spreadsheet program.

21 / 35

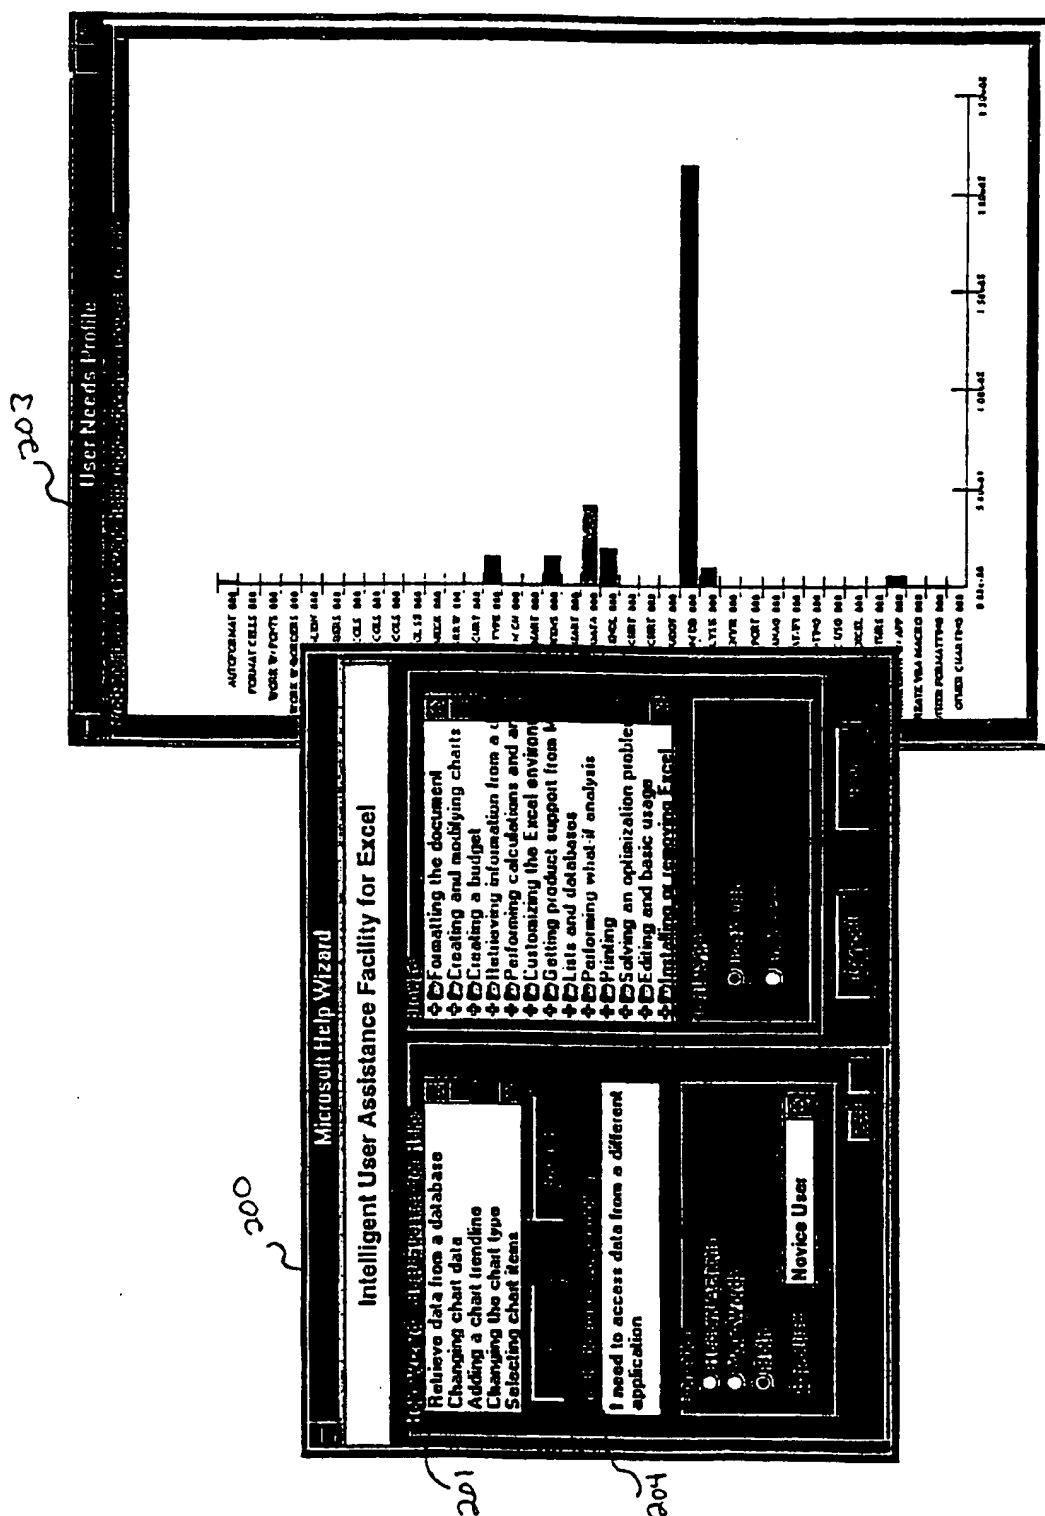


Figure 21. Screen outputs of the IUAF in an exemplary embodiment of a spreadsheet program.

22 / 35




Figure 22. Screen outputs of the IUAF in an exemplary embodiment of a spreadsheet program.




Figure 23. Screen outputs of the IUAF displaying a summary of a custom-tailored tutorial.

24 / 35




Figure 24. Portion of a Bayesian network for intelligent user assistance program.

25 / 35




Figure 25: Portion of Bayesian network for spreadsheet application illustrating modeled (M) and profiled (P) events for charting data of spreadsheet.

26 / 35

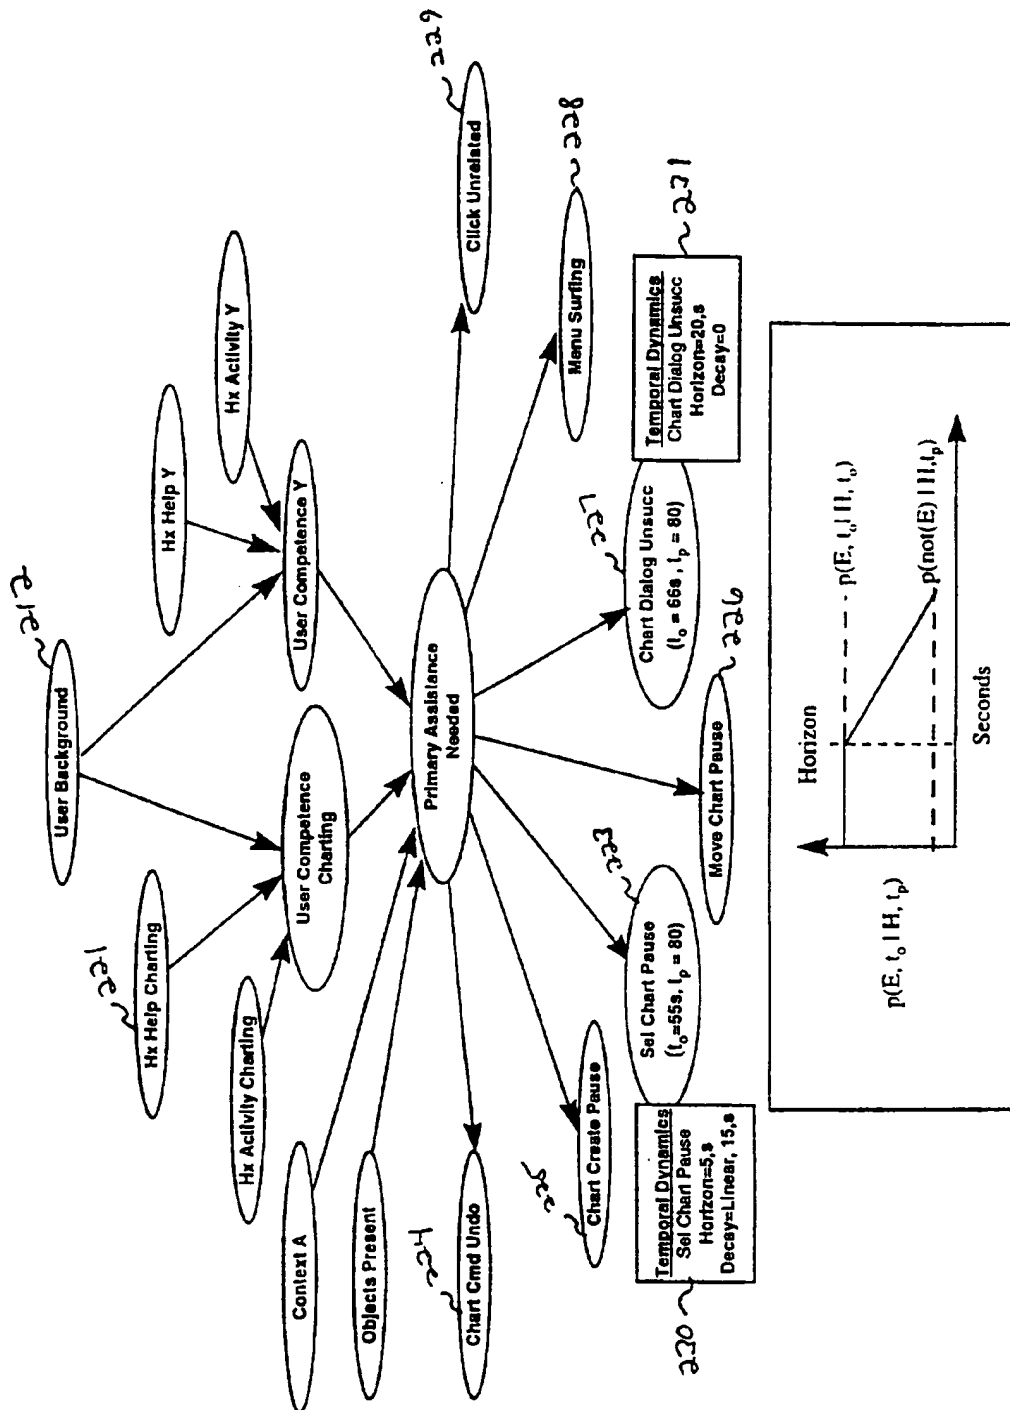


Figure 26: A Bayesian network annotated with temporal dynamics information.

27 / 35




Figure 27. Portion of Bayesian network illustrating model for assisting with exiting in a multiple application setting.




Figure 28. Portion of Bayesian network illustrating reasoning for assisting with exiting in a multiple application setting.




Figure 29. Portion of Bayesian network illustrating reasoning for assisting with exiting in a multiple application setting.




Figure 30. Portion of Bayesian network illustrating reasoning for assisting with exiting in a multiple application setting.

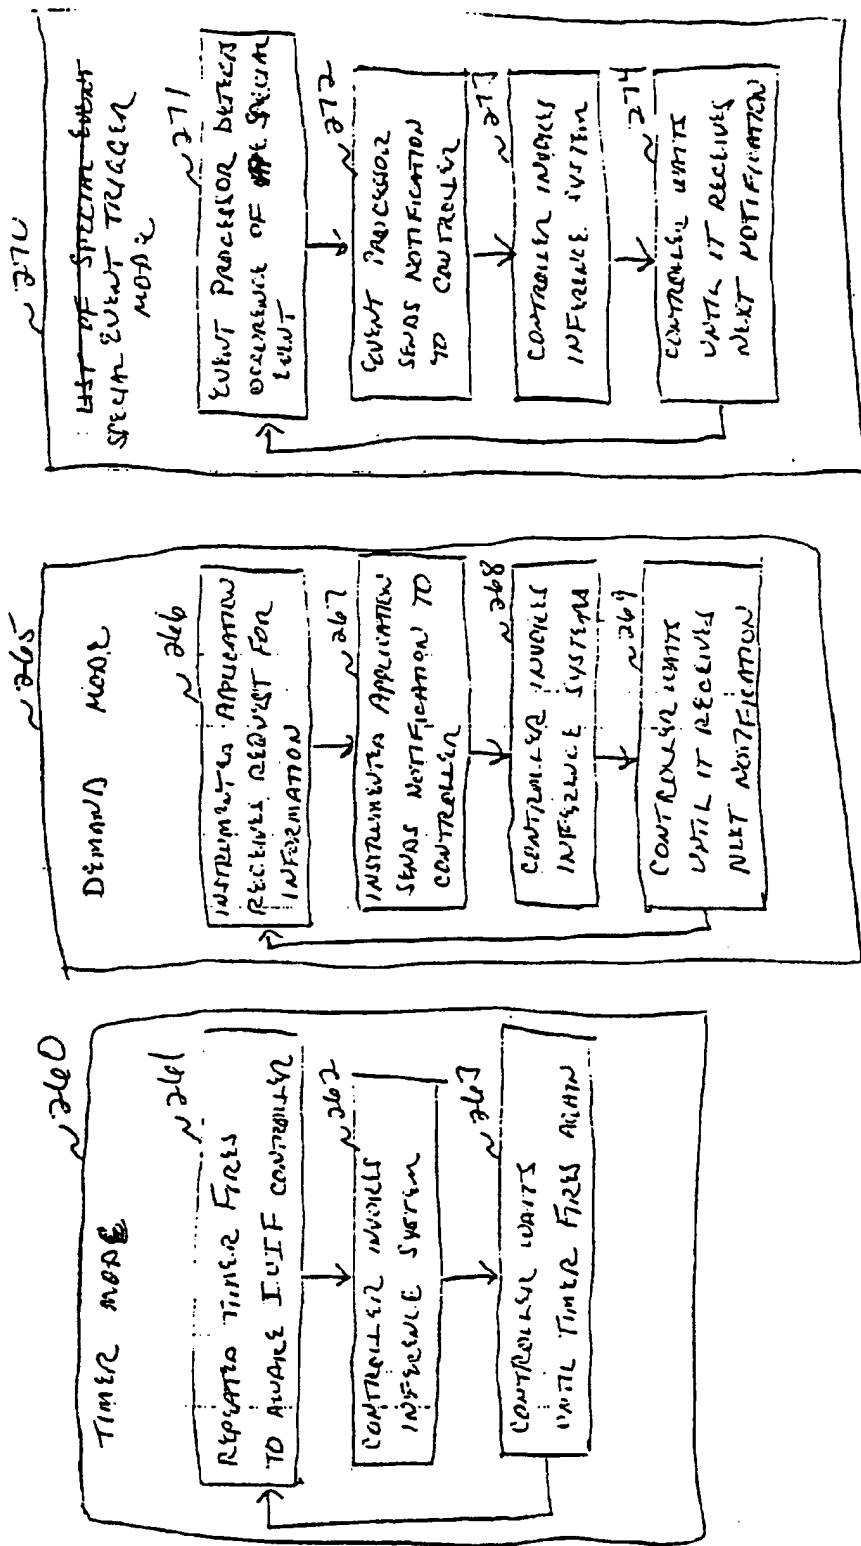


DIAGRAM ILLUSTRATING
CONTROL MODES OF
IUIF CONTROLLER

32/35

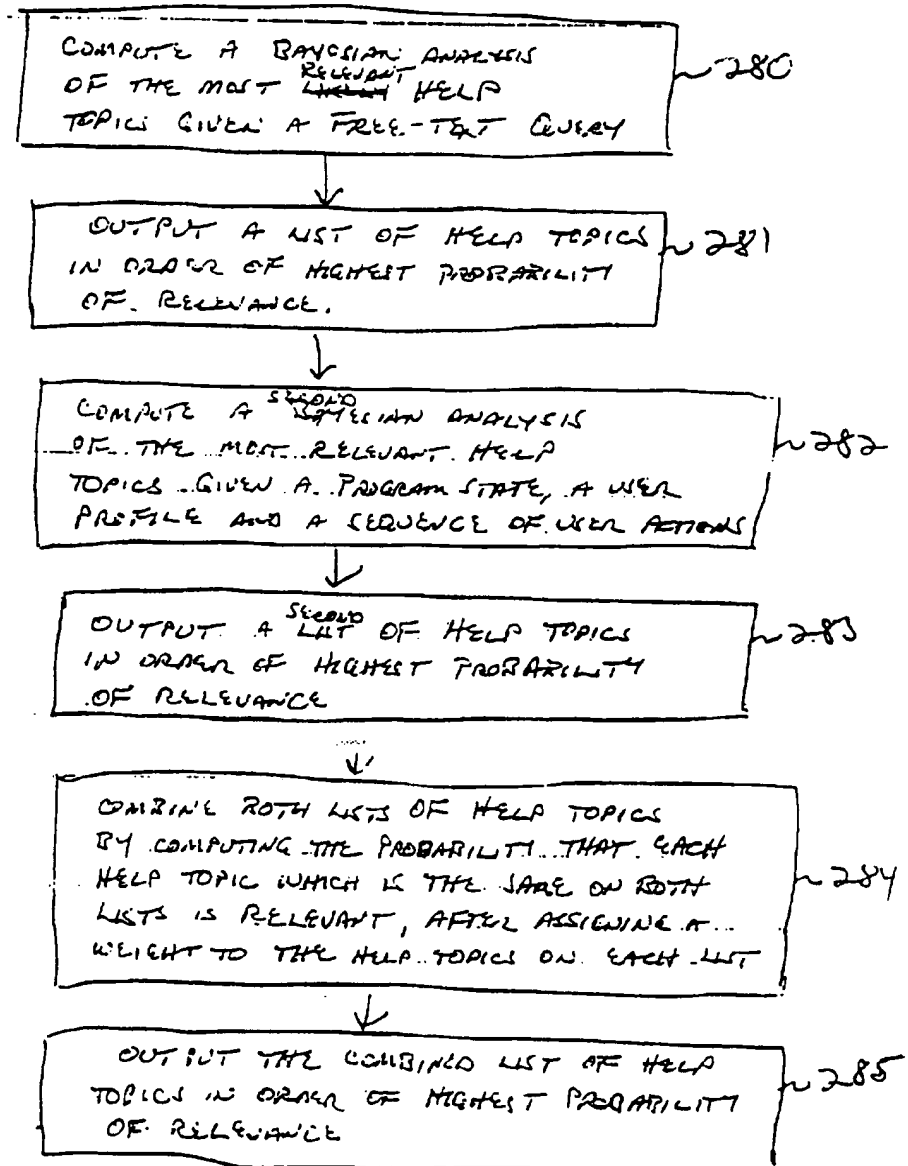


FIG
32

METHOD FOR COMBINING INFERENCE
ANALYSIS OF USER ACTIONS AND PROGRAM
STATE WITH INFERENCE ANALYSIS OF
FREE-TEXT QUERY

33 / 35

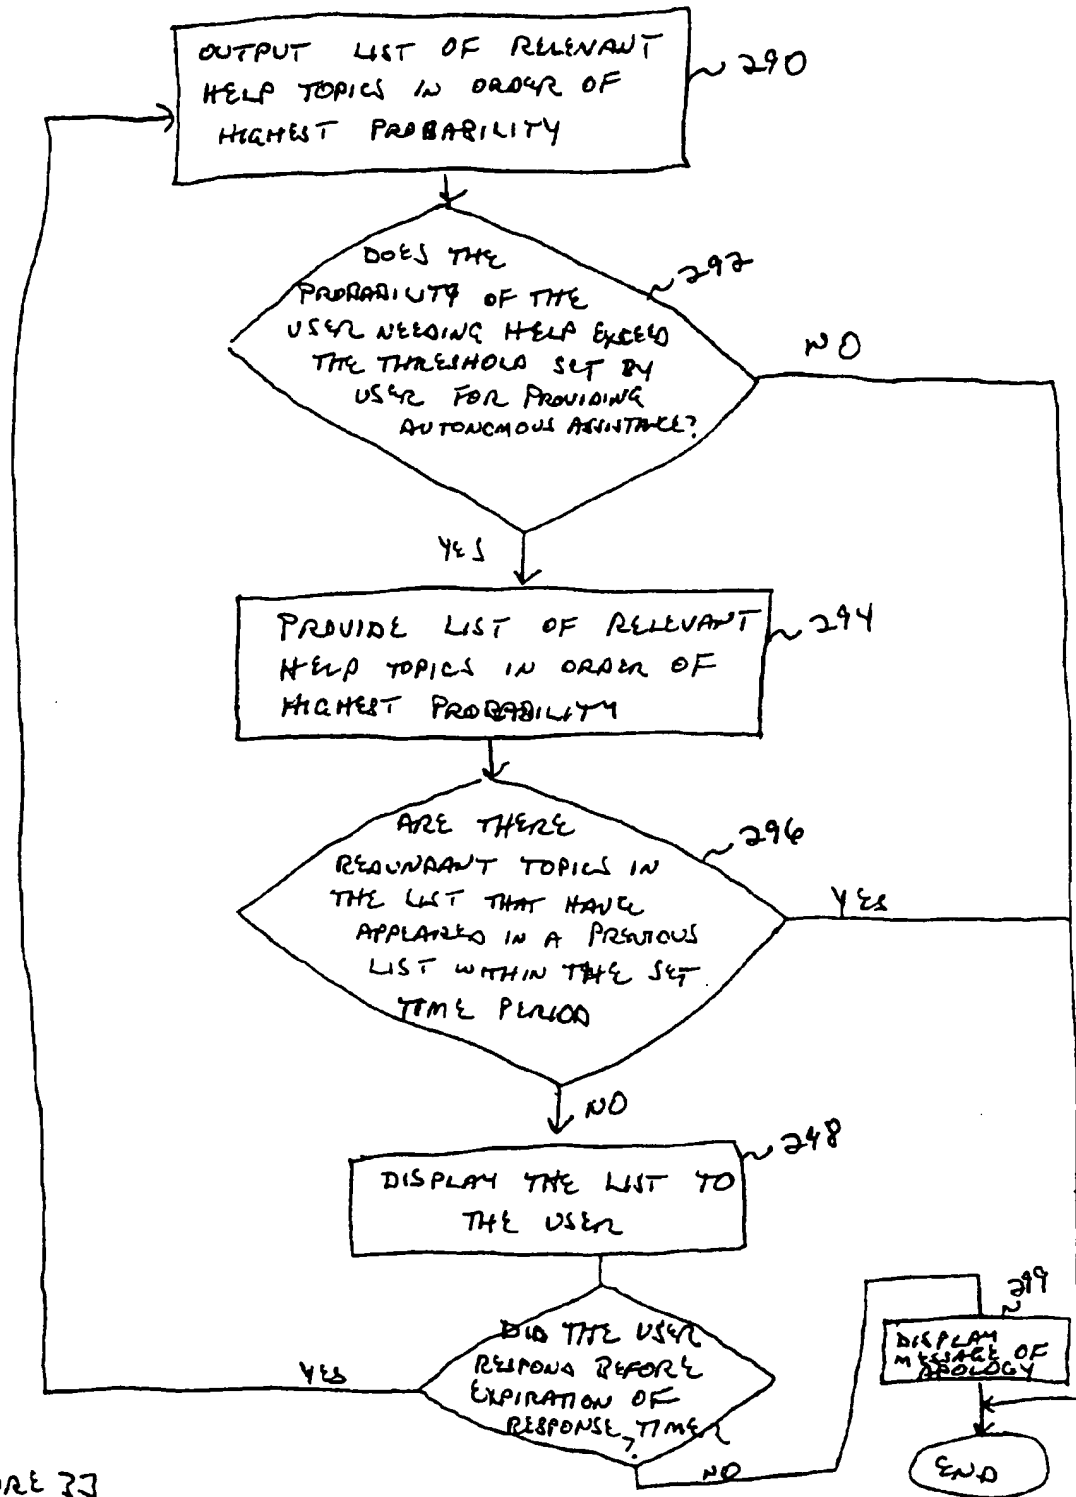


FIGURE 33

34 / 35

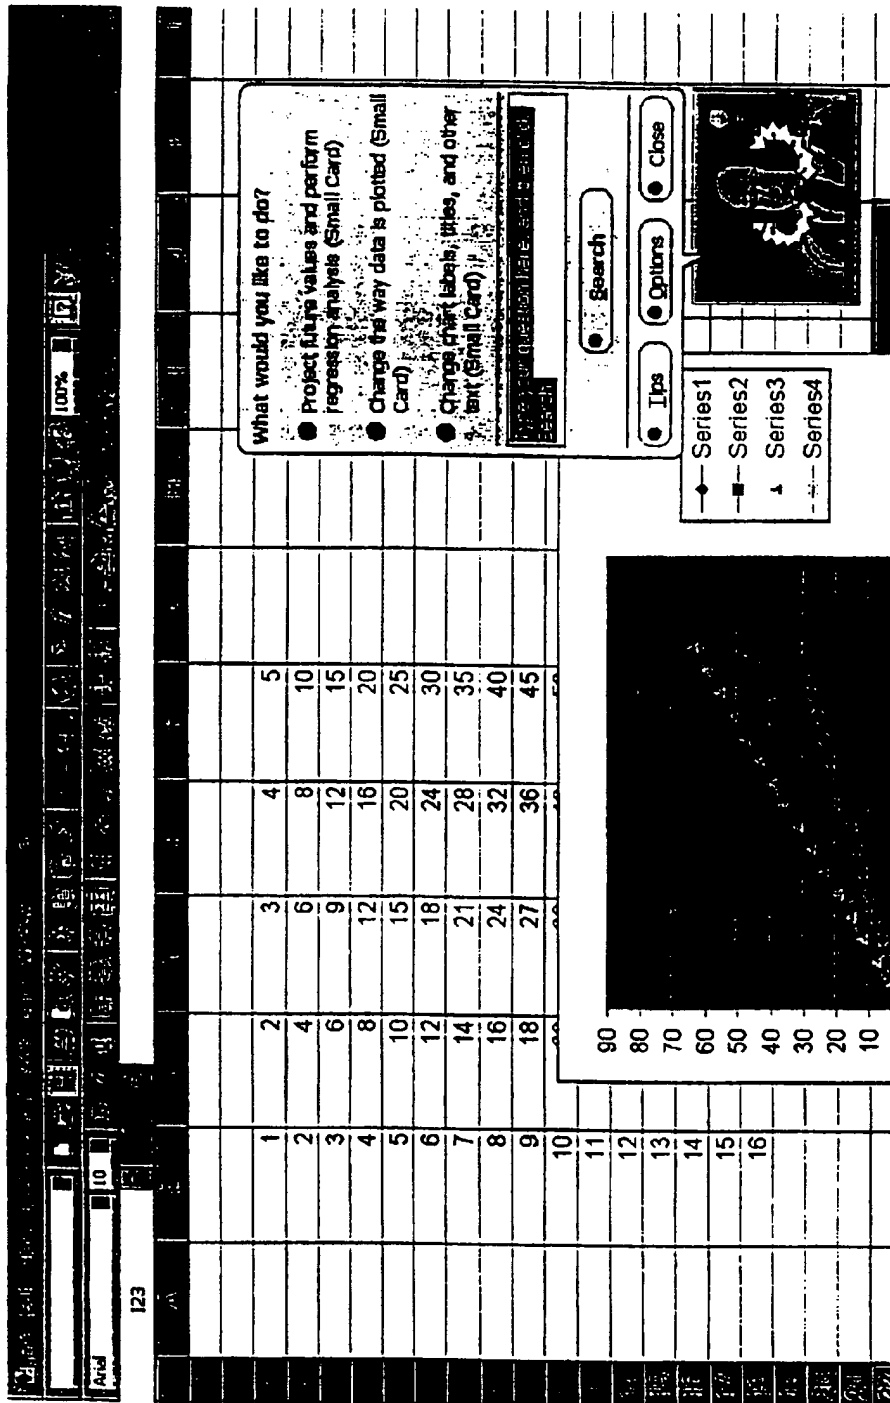


Figure 34

35 / 35

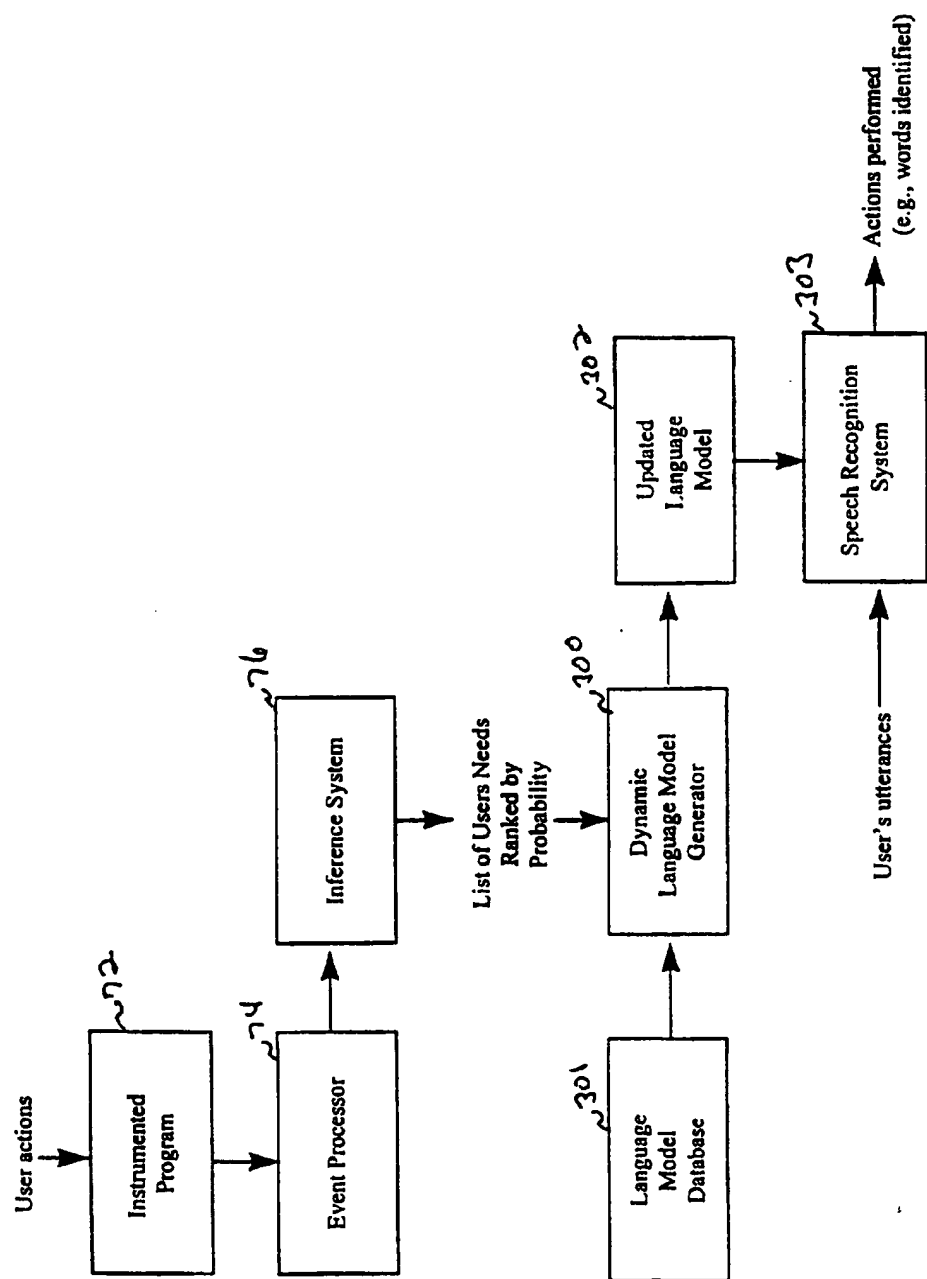


Figure 35. Block diagram of IUAF with speech recognition system showing use of inference results to enhance speech recognition.



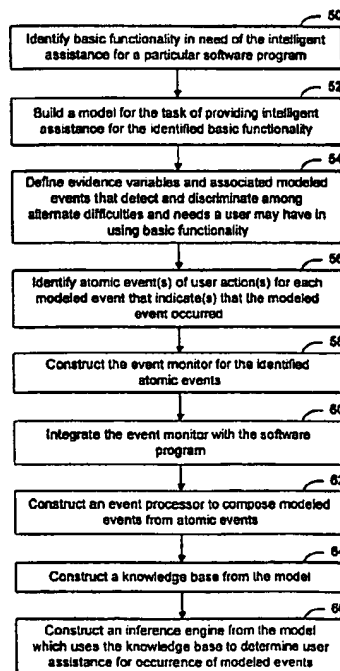
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/44	A3	(11) International Publication Number: WO 98/03907 (43) International Publication Date: 29 January 1998 (29.01.98)
(21) International Application Number: PCT/US97/10856 (22) International Filing Date: 2 July 1997 (02.07.97) (30) Priority Data: 08/684,003 19 July 1996 (19.07.96) US (71) Applicant: MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, WA 98052-6399 (US). (72) Inventors: HORVITZ, Eric; 330 Waverly Way, Kirkland, WA 98033 (US). BREESE, John, S.; 7701 S.E. 58th Street, Mercer Island, WA 98040 (US). HECKERMAN, David, E.; 648 West Lake Sammamish Lane N.E., Bellevue, WA 98008 (US). HOBSON, Samuel, D.; 98 Union Street #905, Seattle, WA 98101 (US). HOVEL, David, O.; 4610 East Lake Sammamish Parkway N.E., Redmond, WA 98053 (US). KLEIN, Adrian, C.; 411 Boylston Avenue E. #203, Seattle, WA 98102 (US). ROMMELSE, Jacobus, A.; Braambos 20, NL-5563 AB Westerhoven (NL). SHAW, Gregory, L.; 5402 Lake Washington Boulevard N.E. #E, Kirkland, WA 98033 (US). (74) Agents: CANNING, Kevin, J. et al.; Seed and Berry LLP, 6300 Columbia Center, 701 Fifth Avenue, Seattle, WA 98104-7092 (US).	(81) Designated States: CN, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i> (88) Date of publication of the international search report: 11 September 1998 (11.09.98)	

(54) Title: INTELLIGENT USER ASSISTANCE FACILITY

(57) Abstract

A general event composing and monitoring system that allows high-level events to be created from combinations of low-level events. An event specification tool allows for rapid development of a general event processor that creates high-level events from combinations of user actions. The event system, in combination with a reasoning system, is able to monitor and perform inference about several classes of events for a variety of purposes. The various classes of events include the current context, the state of key data structures in a program, general sequences of user inputs, including actions with a mouse-controlled cursor while interacting with a graphical user interface, words typed in free-text queries for assistance, visual information about users, such as gaze and gesture information, and speech information. Additionally, a method is provided for building an intelligent user interface system by constructing a reasoning model to compute the probability of alternative user's intentions, goals, or informational needs through analysis of information about a user's actions, program state, and words. The intelligent user interface system monitors user interaction with a software application and applies probabilistic reasoning to sense that the user may need assistance in using a particular feature or to accomplish a specific task. The intelligent user interface also accepts a free-text query from the user asking for help and combines the inference analysis of user actions and program state with an inference analysis of the free-text query. The inference system accesses a rich, updatable user profile system to continually check for competencies and changes assistance that is given based on user competence.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 97/10856

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F9/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	IEEE 1985 PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON CYBERNETICS AND SOCIETY (CAT. NO.85CH2253-3), TUCSON, AZ, USA, 12-15 NOV. 1985, 1985, NEW YORK, NY, USA, IEEE, USA, pages 746-750, XP002051741 BISWAS G ET AL: "An active assistance system for 'vi'"	1-4,9, 19-26, 28-32, 45,47, 49, 52-56, 58,60, 61, 64-67, 79,81, 84,88-95 5,48,50, 51
Y	see the whole document --- -/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

Z document member of the same patent family

Date of the actual completion of the international search

13 January 1998

Date of mailing of the international search report

16. 07.98

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

KINGMA Y.

INTERNATIONAL SEARCH REPORT

Int'l Application No

PCT/US 97/10856

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 92 02880 A (TANDY CORP) 20 February 1992 see abstract; claims 1-3,5,6,8-13; figures 2,3A,3B,6,7 see page 1, line 16 - page 13, line 21 ---	1,3,4, 12,19, 20,27, 49,63, 88,93,94
Y	BAUER M ET AL: "Logic-based plan recognition for intelligent help systems", CURRENT TRENDS IN AI PLANNING. EWSP '93 - 2ND EUROPEAN WORKSHOP ON PLANNING, CURRENT TRENDS IN AI PLANNING. EWSP '93 - 2ND EUROPEAN WORKSHOP ON PLANNING, VADSTENA, SWEDEN, 9-11 DEC. 1993, 1994, AMSTERDAM, NETHERLANDS, IOS PRESS, NETHERLANDS, PAGE(S) 60 - 73 XP002051742 ---	5
X	see page 69, line 1 - page 71, line 37 ---	79,81
X	ICL TECHNICAL JOURNAL, vol. 7, no. 2, 1 November 1990, pages 328-361, XP000176705 SMITH M ET AL: "INTELLIGENT HELP - THE RESULTS OF THE EUROHELP PROJECT" ---	1-4,8,9, 12,19
Y	see abstract ---	13
A	see page 339, line 8 - page 347, line 30 see page 358, line 17 - line 30 ---	22,23
Y	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 31, no. 3, 1 August 1988, XP000106189 "TIME-DEPENDENT UNSOLICITED HELP" see the whole document ---	13
Y	US 5 390 281 A (LUCIW WILLIAM W ET AL) 14 February 1995 see column 10, line 12 - column 11, line 20 see column 12, line 63 - column 16, line 21 -----	48,50,51

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 97/ 10856

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-9, 12, 13, 19-40, 48-52, 54-56, 58, 60-64, 67, 79-88, 92-95, 102-105

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

In. tional Application No

PCT/US 97/10856

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9202880 A	20-02-92	US 5103498 A AU 8401991 A EP 0541712 A	07-04-92 02-03-92 19-05-93
US 5390281 A	14-02-95	US 5477447 A US 5434777 A US 5621903 A US 5644735 A US 5608624 A US 5625814 A	19-12-95 18-07-95 15-04-97 01-07-97 04-03-97 29-04-97

Online Documentation: The Next Generation

JoAnn T. Hackos, PhD
President, Comtech Services, Inc.
710 Kipling Street, Suite 400
Denver, Colorado 80215
303-232-7586
joann.hackos@comtech-serv.com

Information delivered through the computer means many things today. You may embed online help within the software applications you develop, deliver online documentation on CD-ROMs or floppy disks, or establish information databases to be accessed through a company intranet or on the World Wide Web. Some companies distribute nearly all of the information users receive in electronic form rather than as paper manuals. They save the cost of printing and shipping the manuals. Users get the problems of navigating the current crop of online information to use the application software.

The current generation of online information leaves much to be desired from the perspective of the end-user. Certainly, effective online information is out there, designed for quick access and structured so that a user's questions are anticipated. But we find too often that online information is badly designed and inappropriately written.

The most serious problems arise in two ways. First, some organizations simply put their old legacy manuals online. These manuals were designed and written for paper. They work all right if users take the time to print and bind them

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

SIGDOC 97 Snowbird Utah USA

Copyright 1997 ACM 0-89791-861-4/97/10..\$3.50

themselves, or at least they work as well as the old paper manuals worked. The development organization has simply transferred the cost of printing and binding to the customer.

Second, other organizations reconstruct the online help and documentation into WinHelp, Microsoft's tool that links help files to any Windows software application. However, they often reconstruct the online information without sufficient understanding of the types of information users are looking for. In addition, the information may be poorly organized, making navigation difficult. Users searching for information often become frustrated with the lack of appropriate content and the difficulty of finding the information they need.

We can no longer assume that information written for paper can be transferred without change into online forms, nor should we assume that the paper documents were themselves effective. The next generation of online help and other forms of electronic information delivery provides us an opportunity to completely rethink how end-users learn to use application software to do their jobs. Using new design methods, we have the opportunity to provide information that is more useful, less frustrating, less voluminous, and less expensive. If development organizations employ the services of experienced online-information developers and those developers know and use the principles of minimalist design, we can solve many of the problems endemic in today's software documentation and deliver information electronically at the same time.

Using minimalism to develop effective online information

Minimalist design follows the principles originally developed by Dr. John M. Carroll, formerly with IBM's Watson Research Laboratory. Today, those principles are being promoted by other leaders in the field of information design. The principles are simple:

- focus on the user's needs
- ensure ease of access to the right information
- provide just the right words and graphics to ensure successful performance
- emphasize troubleshooting and correcting mistakes
- provide for both novice and expert users through layering

But like any principles founded on a strong understanding of users, minimalist design requires that information designers work closely with users to ensure that their information needs are thoroughly investigated and understood.

Many users tell us that they can't find the information they need in the documentation. In fact, the documentation, whether in paper or online, seems to concentrate more on the intricacies of the product than on answering users' questions. In our research on the usability of online help systems, we often find that unsuccessful documentation focuses on the underlying system rather than the users' needs.

Take, for example, the documentation that came with a cellular phone I bought a few years ago. I travel a great deal, consulting, developing, and teaching workshops on information design. I use a cellular phone to make calls back to the office and to clients, using my long-distance access number. My company bought me a phone that would let me store and recall the 800 number, punch in my access code, and dial the phone number.

Here is the section of the manual that is supposed to explain how to program the phone to dial the access code:

DTMF Tones

During conversation mode, depressing digit keys will send the corresponding DTMF signals

To send a string of DTMF signals while in conversation mode, enter the digits into the display and then press "FCN" + "SND".

For example, if you recall a number programmed with a pause between the telephone number and numbers to be sent as DTMF signals, press "FCN" + "SND" at the pause. The numbers after the pause will be in DTMF format.

Only after talking to a telephone engineer was I able to discover that DTMF tones applied to the task I wanted to perform. Please don't ask what DTMF means—I keep forgetting because it isn't that important to me. To this day, I still can't figure out how to use the DTMF tone function. The instructions appear to be written by an expert and assume that I know all about the inner workings of cellular phones, something I am much too busy with my own work to even want to learn.

Such documentation is typical of an organization that has its development engineers write the manuals. It focuses on the inner workings, the system functions, rather than on what the users want to do. The user seeking procedural help is often not interested in reading information about the underlying structure and logic of the software.

Effective online information focuses on the user

Well-designed and -written online information focuses on the users and what they are trying to do. To develop effective online information, designers must meet the users, watch them doing their jobs, listen to the kinds of questions they ask about the software, and learn how they think about what they are trying to do.

To write effective documentation for the cellular phone, the information designer needs to know that the users want to dial an 800 number and then punch in an access code. Consequently, the designer might write something like the following:

Making credit-card calls automatically

Many credit-card calls require that you enter an access code after you have dialed an 800 number.

To enter an access code automatically, you must store it using the same instructions for storing frequently dialed numbers.

When you want to send your access code, simply recall it from storage and then press **Function Send**.

Since I don't work for the cell phone company, I can't guarantee that these instructions are correct. However, I could follow these instructions because they are written in language I understand.

If we design and write online information with a user focus, rather than a system focus, our users will be able to recognize answers to their questions and points of confusion in the instructional and reference material.

Effective online information focuses on ease of access

Having the right information written in language that the users can understand will go a long way toward improving the quality of the current generation of online documentation. However, the right content is not enough. In many online information systems, the right information is nearly impossible to find. Unless they must look through many volumes of documentation, many people actually find the paper manuals easier to use. Not only can they see what is on your current page, they can see the whole book at one time. By looking at the table of contents, flipping pages, or searching the index, they can usually find the information they need if it is there.

Not so with online information. Unless the information developers have created a consistent and thoughtful structure for the help topics, provided a detailed, expandable table of contents, developed index and full-text searches, and given users a way to flip through topics with a browse sequence, they may never find a topic they need, even if it is there. With most online information available today, users can view only one topic at a time. Users find themselves going through menu after menu of obscure choices before

getting to a topic, only to discover that it doesn't contain the information they need. Sometimes, the only alternative is to start over at the main menu.

Even the search mechanisms in much online information can be difficult to use. If the information has not been thoroughly indexed with lots of synonyms, you may get a series of frustrating "no information found" messages. Even the full-text search may produce hundreds of hits rather than the one relevant topic you are looking for.

Effectively designed online information can be accessible if it is designed with user access in mind. The rule of thumb that I use in my designs reads:

"The user should use no more than three clicks to reach relevant information."

To follow the "three clicks" rule requires that the information designer create a shallow system with a thorough and always-visible table of contents, a thorough index with lots of related words, and a full-text search that leads users through a series of decision points to find what they need.

To aid access to relevant information, the designer must create context-sensitive links to the software application itself. If the user needs information about how to complete a particular field or choose a radio button or check box, that information should be immediately available through some form of "bubble" or "balloon" help. If the user needs information about what to do next, the designer should organize the information according to the most typical workflow. If the user finds graphic information more accessible, the designer should provide a pictorial view of the system with graphic hotspots on the illustration.

Effective online information supports performance

When paper-based information is simply "dumped" online, the user often feels inundated by too many words. Paragraphs that looked readable on paper become unreadable on screen. Information that once occupied two or three

pages now appears in bottomless scrolling text with no end in sight.

Minimalist principles suggest that less information, not more, helps people perform tasks most quickly and easily. Short sentences, step-by-step instructions, brief explanations all help people skim the instructions and get on with the task. They spend less time reading and more time doing.

If users need more than quick job aids, then effectively designed online information presents more information through hypertext links and layering. As the first line of information, the users' needs may be addressed through "cue cards" that sit on top of the appropriate application screen. Such cue cards can even assist the users in performing the task by interacting directly with the software. If the designer includes interactive help screens, the users can select an icon or push a button that performs part of the task for them.

Effective online information emphasizes troubleshooting

Users tell us that the information they most need is information that helps them correct mistakes and get out of trouble. In many cases, something has happened in the software that the users don't understand. They would like to ask "What happened?" and have the online help respond appropriately.

Well-designed online information answers the critical question "Why?" Ideally, the designers have identified the most likely problems that will occur for ordinary users and provided solutions and troubleshooting tips.

Effective online information clusters information for different user types

Users differ in their information needs. Experts are looking for help performing more complex tasks, while beginners want help with the basics. Experts want in-depth explanations of how the software works, while beginners only want to complete a task as quickly as possible.

When information is well structured for a range of user types, then users can move from simple step-by-step instructions to detailed background information to examples and demonstrations,

even to overviews and discussions of concepts. Tools like Windows Help and other hypertext-based systems often allow designers to create layers of information around a central help topic. Users can access pop-up definitions of terms without leaving their help topic. They can get more help to perform difficult steps in a procedure by linking to an expanded step. Experts can click on buttons that bring up in-depth explanations; beginners can ask for demos or examples of how a particular function will meet their needs.

A rich information architecture based on clusters of information can enhance the users' options and allow users to perform tasks more effectively.

Minimalism summarized

The injunctions of a minimalist approach focus on the need to plan help that is genuinely helpful. To do so is not as simple as it sounds. Good information design begins with the design of a user-centered interface. Information designers recognize that the interface contains information that enables people to learn and perform actions and reach their goals. But the information does not stop with the interface. As human-performance specialists, we know that users have different information needs, that novices do not have the same needs or interpret information in the same way as experts.

Types of online information

The design of effective online information starts with the interface and then moves on to context-sensitive help; procedural, conceptual, reference, and instructional information; interactive multimedia tutorials; and more. These online information types are described in the next section.

Beginning with the user interface

The next generation of online documentation will begin with the user interface. The interface provides information directly to users through text, graphics, and icons. It also provides information indirectly, through the arrangement of the screen objects and the workflow encouraged by the sequence of events supported. The interface is, however, just the beginning. More extensive online information is already being included as part of the interface in the form of wizards, coaches, and guides. Apple Guide, for

example, provides users the ability to work within the help system to accomplish tasks. The wizards included in parts of Microsoft Office provide step-by-step instructions inside the application software. The line that many developers perceive between the interface and help has becoming increasingly fuzzy.

Adding bubble help

The institution of bubble or balloon help has enabled information designers to include a second level of information to support user tasks. If the name of a function is unclear, if the icons are indecipherable, then the user can request help that answers questions such as

"What is this?"

"What do I do with this?"

"What information do I enter here?"

The questions are answered in the context of the functions; the information is immediately at hand, rather than buried in a paper user manual or somewhere in the online manual. Many users have come to view the immediate help available on the screen as simply an extension of the interface, another level of information.

Creating context-sensitive procedural help

The next level of information comes with context-sensitive help. Although bubble help is itself context-sensitive, we generally refer to more extensive step-by-step instruction as providing the core of context-sensitive help. This type of help addresses the critical user questions of

"How do I . . . ?"

In the current design of most help systems, users must initiate the call for help. What they most often expect to receive in response is procedural help, explaining how to complete tasks and reach relevant goals. For example, if users want to set the tabs in the cells of a table, they should be able to find step-by-step procedures for performing this task in the online help.

With context-sensitive procedural help, the user can find appropriate step-by-step procedures for the task at hand because the help designer has anticipated the user's questions. The designer has

looked at the content of the interface and decided what type of help might be needed by someone finding himself in this part of the software. The success of context-sensitive help depends heavily on the ability of the designer to anticipate the user's need for information in a particular context.

A very few help systems attempt to trace the keystrokes that the user has made in an effort to better anticipate the help needed. In general, however, such systems have been unsuccessful because it has proven extraordinarily difficult to relate a series of keystrokes to a user problem.

Procedural help can also be accessed by the user through a variety of mechanisms provided in most help systems today, including tables of contents, full-text search, index, and other implementations of keyword searches such as Microsoft's Answer Wizard in Windows 95. Once again, the user initiates the search for procedural help by searching through index entries or tables of contents or by manipulating a request through a search system.

Including conceptual, background, and reference information

In addition to procedural information, users may be interested in reading conceptual, background, and reference information, most of which can be provided in some electronic form. Reference information is especially well suited for online use. The information describing UNIX commands, for example, has long been available online through the MAN pages.

Conceptual information is somewhat less successful on line. Many users prefer to read conceptual information carefully and at leisure. They want to be able to take home a book about the problem they are trying to solve. The book, however, can be distributed electronically, available for on-demand printing and updating through a web site.

We also have the opportunity with fully integrated information to provide what has been often thought more traditionally regarded as training. Online tutorials, interactive multimedia training, and web-based training can be included in the context of the software interface. One system, for example, provides users with a help

icon to access the help system and with a tutor icon to access the computer-aided instruction (CAI). Demonstrations and interactive tutorials can also be embedded within the help system itself. In addition to accessing procedures and conceptual text, users can decide to try an interactive example. Apple provides such a system in the Macintosh Guide, which will actually perform a step in a procedure, such as selecting a menu item or opening a dialog box if the users cannot figure out the step for themselves. In Windows 95™, users have the option of asking the help system to locate and open a dialog box.

It is also possible to associate a help system with an online training simulation. In one application, we used the actual context-sensitive help from inside a CAI simulation. In other designs, we have layered information so that the users have many choices depending on their information needs. From a procedural help topic, they might select a pop-up definition, a conceptual overview, an annotated example, an interactive tutorial, or many of the myriad possibilities available to the information designer.

Moving into the next generation

All of these methods of providing help and information to the users from the interface through interactive multimedia add up to a system that supports the performance of users from novices to experts. Performance-support systems integrate many types and levels of information into the interface, allowing users to find the level of support that best suits their needs. In fact, much of the discussion of performance support in the second half of the 90s is about the design of user-centered supportive interfaces and functionality.

Today, we are extending information design and performance support from the application and the CD-ROM to the Internet and the intranet. Both Netscape and Microsoft have begun to provide functionality that allows designers to link applications to web pages. The web pages can be designed to provide procedural, conceptual, reference, and instructional help. They might also be designed to enable links to people-based customer support. Customer support is, in turn,

supported by electronically based information constructed in databases (knowledge bases) and around frequently asked questions (FAQs).

What should we look for in information design in the next generation, over the next five to ten years? I hope that we will see increasingly sophisticated interface designs that support user workflow and are more successful in anticipating user problems and questions. In a recent interview with Austin Henderson, which appeared in *Common Ground*, the quarterly publication of the Usability Professionals' Association, I described Henderson's work at Xerox PARC (Palo Alto Research Center) and now as manager of the Discourse Architecture Lab, part of Apple Research Laboratories. He discussed the need for developers to be much smarter about accommodating application design to users. He would like to see more anticipation of user problems and the design of systems that can initiate help by creating a dialog with the users. Such a dialog helps create a more custom response to individual users or users working in the context of an organization and its requirements.

Many of us in user-centered interface design have predicted for some time now the demise of large, multi-functional software tools in favor of more specifically designed, customized applications. We are finally seeing this prediction come to fruition with the development of applets for the internet, which are small applications used for very specific purposes. The applets can be customized to the needs of users in specialized work environments. To develop such focused applications will, however, require considerable understanding of the users, tasks, and environment. The information needed by users to understand what they are doing and how to use the application will come from their environment and the information built into the interface and help.

The next generation of online information will depend, in part, on these advances in technology. However, technological advances will not improve usability or help us develop strong design methods. For online information to succeed will require superb design—or the information we disseminate electronically will simply add to the information glut.

4/9,K/1 (Item 1 from file: 4)

DIALOG(R)File 4:INSPEC

(c) 2004 Institution of Electrical Engineers. All rts. reserv.

6625486 INSPEC Abstract Number: C2000-08-6150E-002

Title: Does help really help? [Terminology-related problems with online help systems]

Author(s): Schweibenz, W.

Journal: NFD Information - Wissenschaft und Praxis vol.51, no.4 p.

199-207

Publisher: Deutsche Gesellschaft fur Dok,

Publication Date: June 2000 Country of Publication: Germany

CODEN: NADOAW ISSN: 1434-4653

SICI: 1434-4653(200006)51:4L;1-P

Material Identity Number: G396-2000-005

Language: German Document Type: Journal Paper (JP)

Treatment: Practical (P); Experimental (X)

Abstract: Terminology-related problems affect the usability of software documentation both in print and electronic form. Terminology causes difficulties in accessing help information and in applying help information. These problems reduce the usability of software documentation. Agents like the Answer Wizard could be helpful and enhance the usability of software documentation. A usability study with thinking aloud examines how terminology-related problems affect the use of the Microsoft Answer Wizard "Clip-it for Microsoft Office 97" and how they reduce the usability of the online documentation. (0 Refs)

Subfile: C

Descriptors: help systems; human factors; nomenclature; office automation ; software agents; system documentation

Identifiers: online help systems; terminology-related problems; software documentation usability; information access; help agents; Microsoft Answer Wizard; thinking aloud; Clip-it for Microsoft Office 97; online documentation

Class Codes: C6150E (General utility programs); C7104 (Office automation); C0240 (Ergonomic aspects of computing)

Copyright 2000, IEE

...Abstract: in applying help information. These problems reduce the usability of software documentation. Agents like the Answer Wizard could be helpful and enhance the usability of software documentation. A usability study with thinking aloud examines how terminology-related problems affect the use of the Microsoft Answer Wizard "Clip-it for Microsoft Office 97" and how they reduce the usability of the online documentation.

...Identifiers: Microsoft Answer Wizard ; ...

...Clip-it for Microsoft Office 97

?